

MSC-900 Network Configuration Manual

MANU0222-01 - Rev. A - June 18, 1997

Software Version 1.0.x

FORE Systems, Inc.

1000 FORE Drive Warrendale, PA 15086-7502 Phone: 412-742-4444

FAX: 412-742-7742

http://www.fore.com

Legal Notices

Copyright [©] 1995-1997 FORE Systems, Inc. All rights reserved. FORE Systems is a registered trademark, and *ForeRunner*, *ForeRunnerLE*, *ForeThought*, *ForeView*, *PowerHub*, and *CellPath* are trademarks of FORE Systems, Inc. All other brands or product names are trademarks or registered trademarks of their respective holders.

U.S. Government Restricted Rights. If you are licensing the Software on behalf of the U.S. Government ("Government"), the following provisions apply to you. If the Software is supplied to the Department of Defense ("DoD"), it is classified as "Commercial Computer Software" under paragraph 252.227-7014 of the DoD Supplement to the Federal Acquisition Regulations ("DFARS") (or any successor regulations) and the Government is acquiring only the license rights granted herein (the license rights customarily provided to non-Government users). If the Software is supplied to any unit or agency of the Government other than DoD, it is classified as "Restricted Computer Software" and the Government's rights in the Software are defined in paragraph 52.227-19 of the Federal Acquisition Regulations ("FAR") (or any successor regulations) or, in the cases of NASA, in paragraph 18.52.227-86 of the NASA Supplement to the FAR (or any successor regulations).

Printed in the USA.

No part of this work covered by copyright may be reproduced in any form. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

This publication is provided by FORE Systems, Inc. "as-is" without warranty of any kind, either express or implied, including, but not limited to, the implied warranties or conditions of merchantability or fitness for a particular purpose. FORE Systems, Inc. shall not be liable for any errors or omissions which may occur in this publication, nor for incidental or consequential damages of any kind resulting from the furnishing, performance, or use of this publication.

Information published here is current or planned as of the date of publication of this document. Because we are improving and adding features to our products continuously, the information in this document is subject to change without notice.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 (October 1988) and FAR 52.227-19 (June 1987).

The VxWorks software used in the Mini Loader is licensed from Wind River Systems, Inc., Copyright ©1984-1996.

FCC CLASS A NOTICE

WARNING: Changes or modifications to this unit not expressly approved by the party responsible for compliance could void this user's authority to operate this equipment.

NOTE: The MSC-900 has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of the equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

DOC CLASS A NOTICE

This digital apparatus does not exceed Class A limits for radio noise emission for a digital device as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le present appareil numerique n'emet pas de bruits radioelectriques depassant les limites applicables aux appareils numeriques de la class A prescrites dans le reglement sur le brouillage radioelectrique edicte par le ministere des Communications du Canada.

VCCI CLASS 1 NOTICE

この装置は、第一種情報処理装置(商工業地域において使用されるべき情報処理装置)で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会(VCCI)基準に適合しております。

従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

This equipment is in the Class 1 category (Information Technology Equipment to be used in commercial and/or industrial areas) and conforms to the standards set by the Voluntary Control Council For Interference by Information Technology Equipment aimed at preventing radio interference in commercial and/or industrial areas. Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers, etc. Read the instructions for correct handling.

FCC REQUIREMENTS (Notice to Users of DS1 Service)

The following instructions are provided to ensure compliance with the Federal Communications Commission (FCC) Rules, Part 68.

- This device must only be connected to the DS1 network connected behind an FCC Part 68 registered channel service unit. Direct connection is not allowed.
- (2) Before connecting your unit, you must inform the telephone company of the following information:

Port ID	REN/SOC	FIC	USOC
NM-6/DS1C	6.0N	04DU9-BN,	RJ48C
		04DU9-DN,	
NM-2/DS1C	6.0N	04DU9-1ZN, and	RJ48C
		04DU9-1SN	

- (3) If the unit appears to be malfunctioning, it should be disconnected from the telephone lines until you learn if your equipment or the telephone line is the source of the trouble. If your equipment needs repair, it should not be reconnected until it is repaired.
- (4) If the telephone company finds that this equipment is exceeding tolerable parameters, the telephone company can temporarily disconnect service, although they will attempt to give you advance notice if possible.
- (5) Under the FCC Rules, no customer is authorized to repair this equipment. This restriction applies regardless of whether the equipment is in or out of warranty.
- (6) If the telephone company alters their equipment in a manner that will affect use of this device, they must give you advance warning so as to give you the opportunity for uninterrupted service. You will be advised of your right to file a complaint with the FCC.

CANADIAN IC CS-03 COMPLIANCE STATEMENT

NOTICE: The Industry Canada label identifies certified equipment. This certification means that the equipment meets certain telecommunications network protective, operational and safety requirements. The Industry Canada label does not guarantee the equipment will operate to the user's satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local tele-communications company. The equipment must also be installed using an acceptable method of connection. In some cases, the company's inside wiring associated with a single line individual service may be extended by means of a certified connector assembly (telephone extension cord). The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations.

Repairs to certified equipment should be made by an authorized Canadian maintenance facility designated by the supplier. Any repairs or alterations made by the user to this equipment, or equipment malfunctions, may give the telecommunications company cause to request the user to disconnect the equipment.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines and internal metallic water pipe system, if present, are connected together. This precaution may be particularly important in rural areas.

<u>Caution</u>: Users should not attempt to make such connections themselves, but should contact the appropriate electric inspection authority, or electrician, as appropriate.

E1 AND E3 NOTICE

The E1 (NM-6/E1C and NM-2/E1C) and E3 (NM-4/E3C and NM-2/E3C) network modules that are described in this manual are approved for use in FORE Systems' host systems providing that the instructions below are strictly observed. Failure to follow these instructions invalidates the approval.

Pan European Approval - CE Marking

Pan European approval of the E1 network module was issued by BABT following assessment against CTR12. This means that it can be connected to ONP and unstructured PTO-provided private circuits with 120 Ω interfaces in all European countries, according to Telecommunications Terminal Equipment (TTE) Directive 91/263/EEC. Thus, the following CE mark applies:

C€168.X

The E1 and E3 network modules conform to safety standard EN60950 1992 following the provisions of Low Voltage Product Safety Directive 73/23/EEC and CE Marking Directive 93/68/EEC, and can be marked accordingly with the CE symbol.

The E1 and E3 network modules conform to EN55022 1994 and EN50082-1 1992 following the provisions of the EMC Directive 89/336/EEC, and can be marked accordingly with the CE symbol.

National Approvals

UK

Network Module	Connects to	Approval Number
E1	Structured and unstructured PTO-provided private circuits with 75 Ω interfaces	AA60953
E3	PTO-provided private circuits with 75 Ω interfaces	NS/4387/1/T/605954

Germany

Network Module	Connects to	Approval Number
E3	Structured PTO-provided private circuits with 75 Ω interfaces	A127535H for the MSC-900

Switzerland

Network Module	Connects to	Approval Number
E1	Structured PTO-provided private circuits with 120 Ω interfaces	96.0872.J.N
E3	Structured PTO-provided private circuits with 75 Ω interfaces	96.0873.J.N

Required User Guide Statements - UK Installation

The use of auxiliary products not authorized by FORE Systems in FORE Systems ATM Switches may cause the power specification to be exceeded and is a potential safety hazard.

The equipment must be installed such that with the exception of the connections to the host, clearance and creepage distances shown in the table below are maintained between the network module and any other assemblies which use or generate a voltage shown in the table below. The larger distance shown in brackets applies where the local environment within the host is subject to conductive pollution or dry non-conductive pollution which could become conductive due to condensation. Failure to maintain these minimum distances invalidates the approval.

Clearance (mm)	Creepage (mm)	Voltage Used or Generated by Host or by Network Modules
2.0	2.4 (3.8)	Up to 50 V _{rms} or V _{dc}
2.6	3.0 (4.8)	Up to 125 V _{rms} or V _{dc}
4.0	5.0 (8.0)	Up to 250 V _{rms} or V _{dc}
4.6	6.4 (10.0)	Up to 300 V _{rms} or V _{dc}

For a host or other expansion card fitted in the host, using or generating voltages greater	Above 300 V_{rms} or V_{dc}
than 300V (rms or dc), advice from a competent telecommunications engineer must be	
obtained before installation of the relevant equipment.	

NOTE: Installing the network modules in the appropriate FORE Systems hosts, according to the installation instructions provided, satisfies the requirements listed above.

The following tables show the available ports and their safety status:

NM-6/E1C and NM-2/E1C

Ports	Safety Status
E1 Ports	TNV operating at SELV
Bus Connector	SELV

NM-4/E3C and NM-2/E3C

Ports	Safety Status
E3 Ports	TNV operating at SELV
Bus Connector	SELV

CE NOTICE

Marking by the symbol **CE** indicates compliance of this system to the EMC (Electromagnetic Compatibility) directive of the European Community and compliance to the Low Voltage (Safety) Directive. Such marking is indicative that this system meets or exceeds the following technical standards:

- EN 55022 "Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment."
- EN 50082-1 "Electromagnetic compatibility Generic immunity standard Part 1: Residential, commercial, and light industry."
- IEC 1000-4-2 "Electromagnetic compatibility for industrial-process measurement and control equipment Part 2: Electrostatic discharge requirements."
- IEC 1000-4-3 "Electromagnetic compatibility for industrial-process measurement and control equipment Part 3: Radiate electromagnetic field requirements."
- IEC 1000-4-4 "Electromagnetic compatibility for industrial-process measurement and control equipment Part 4: Electrical fast transient/burst requirements."

SAFETY CERTIFICATIONS

ETL certified to meet Information Technology Equipment safety standards UL 1950, CSA 22.2 No. 950, and EN 60950.

Table of Contents

Preface

	-	
Techni Typogi Import Laser	ical Suppo raphical S tant Inform Notice Precaution Modifica Placeme	aries. i bort iii tyles iii nation Indicators iv .v .v ons. vi ations to Equipment vi ent of a FORE Systems Product vi Cord Connection vii
СНАР	TER 1	Configuring PVCs
1.1	General	Concepts1 - 1
1.2	Virtual F	Paths1 - 3
	1.2.1 1.2.2	Through Paths
1.3		/irtual Paths
1.0	1.3.1 1.3.2	Listing Through Paths
1.4	Virtual (Channels
	1.4.1 1.4.2	Soft Permanent Virtual Circuits
1.5	Creating	g PVCs and SPVCs

1.5.1

1.5.2

1.5.3

1.5.4

1.5.1.1

1.5.1.2

1.5.2.1

1.5.2.2 1.5.2.3

1.5.4.1

Table of Contents

1.6		ypes
1.7	Traffic P	olicing (Usage Parameter Control)
	1.7.1	Leaky Bucket Algorithm1 - 32
	1.7.2	Non-conforming Cells: Tagging vs. Dropping
	1.7.3	UNI 3.0 UPC Traffic Contract Parameters
	1.7.4	AMI UPC Commands
CHAP	TER 2	Configuring a Classical IP ATM Network
2.1	Introduc	tion
	2.1.1	Logical IP Subnets
	2.1.2	Classical IP Interfaces
	2.1.3	SPANS Interface 2 - 3
2.2		Registration and ILMI
	2.2.1	NSAP Addresses2 - 4
	2.2.2	Operating with ILMI Support
	2.2.3 2.2.4	Operating without ILMI Support
2.3		Configuration
2.3		
	2.3.1 2.3.2	Theory
		2.3.2.1 Configuring a Host as an ARP Server for One-Time Use 2 - 8
		2.3.2.2 Configuring a Host as a Permanent ARP Server
		2.3.2.3 Configuring a Hewlett-Packard Computer as an ARP Server . 2 - 9
	2.3.3	Configuring an MSC-900 to be an ARP Server
	2.3.4	Classical IP Operation
	2.3.5	Operational Issues
2.4		al IP PVCs
	2.4.1	Theory and Configuration
	2.4.2	Revalidation and Removal
2.5		ing
2.6	_	ring the Network
	2.6.1	Third-Party Host with No ILMI and No RFC-1577 Support 2 - 16
	2.6.2	Third-Party Switch with ILMI and No RFC-1577 Support 2 - 17
	2.6.3	Third-Party Switch with RFC-1577 and No ILMI Support 2 - 18
CHAP	TER 3	Configuring an Emulated LAN
3.1	Introduc	tion
	3.1.1	Ethernet ELANs
	3.1.2	Token Ring ELANs 3 - 1

3.2	Emula	ed LAN Components.		- 2
	3.2.1	LAN Emulation Clie	ent (LEC)	- 3
	3.2.2		nfiguration Server (LECS)	
	3.2.3		ver (LES)	
	3.2.4		nown Server (BUS)	
3.3	Emula	ed LAN Operation		- 4
	3.3.1	Initialization		- 6
	3.3.2	Registration and Ac	ddress Resolution	- 6
	3.3.3	Data Transfer		- 7
3.4	Config	ring an ELAN		- 8
	3.4.1		CS Configuration Database File	
			Begin3	
			iguration File Syntax	
			n ELAN	
			Client	
			rol Parameters	
	3.4.2		figuration File	
	3.4.3		Configuration File	
	3.4.4		mulation Services	
			e LECS	
	.		LES and a BUS	
	3.4.5		and Joining an ELAN	
			LEC	
		3.4.5.2 Configuring	g the LEC Failover Mechanism	26
CHAP	TER 4	Configuring Circuit	Emulation Services	
4.1	Config	ring CES Connections	s	- 2
	4.1.1	Creating a New CE	S Connection	- 2
	4.1.2		nnections	
APPE	NDIX A	Configuring SNMP		
A.1	SNMP	ndexing	A	- 1
A.2	SNMP	Traps	A	- 3
	A.2.1	•	Destinations	
	A.2.2		rap DestinationsA -	
	A.2.3	. , ,	rap Destinations	

Table of Contents

APPEN	IDIX B F	oreThought PNNI
B.1	FT-PNNI I	Routing
	B.1.1	Hello Protocol
	B.1.2	Topology Database Exchange B - 2
	B.1.3	Flooding
	B.1.4	Hierarchical Routing
	В	.1.4.1 Hierarchical Addressing
		B.1.4.1.1 Switch Prefix
		B.1.4.1.2 Switch Summary Prefix
	р	B.1.4.1.3 Peer Group ID
D 0		.1.4.2 Path Computation
B.2	•	ical Network
	B.2.1	Peer Groups
	B.2.2	Peer Group Topology
	B.2.3 B.2.4	Border Switches
	B.2.4 B.2.5	Peer Group Summary Node (PGSN)
	B.2.6	Single Switch Perspective
APPEN	IDIX C U	INI Signalling
C.1	VCI Alloca	ation Range
	C.1.1	Determining the VCI Allocation Range with ILMI Down C - 2
	C.1.2	Determining the VCI Allocation Range with ILMI Up
C.2	Non-path	Associated Signalling Enhancements
	C.2.1	VC-Space
	C.2.2	Dynamic Paths
C.3	UNI Auto	Configuration Procedures
C.4	Allowable	Combination of Traffic Parameters
Acrony	ms	
Glossa	ry	
Index		

Preface

This manual provides the technical information needed to configure FORE Systems' MSC-900, it's supported FORE Systems network modules, and the accompanying *ForeThought-SP* software. This document also provides general ATM information and general product information. This document was created for users with various levels of experience. If you have any questions or problems, please contact FORE Systems' Technical Support.

Chapter Summaries

Chapter 1 - Configuring PVCs - Describes how to create PVCs through the ATM Management Interface (AMI).

Chapter 2 - Configuring a Classical IP ATM Network - Describes how to design, configure, and maintain a Classical IP ATM network.

Chapter 3 - Configuring an Emulated LAN - Provides an overview of LAN Emulation and gives an example of how to configure an Emulated LAN.

Chapter 4 - Configuring Circuit Emulation Services - Contains configuration information for supporting Circuit Emulation (CEM) network modules.

Appendix A - Configuring SNMP - Describes remote SNMP configuration of the MSC-900.

Appendix B - *ForeThought* **PNNI** - Describes how this scalable routing and signalling protocol can be used to simplify large network topologies.

Appendix C - UNI Signalling - Describes UNI signalling in *ForeThought*-SP 1.0.x, including VCI allocation range, signalling scope enhancements, UNI auto-configuration, and traffic parameters.

Acronyms - Contains a list of common networking acronyms.

Glossary - Contains definitions for networking terms, both general and ATM-specific.

Technical Support

In the U.S.A., you can contact FORE Systems' Technical Support by any one of four methods:

1. If you have access to Internet, you may contact FORE Systems' Technical Support via e-mail at:

support@fore.com

2. You may telephone your questions to "support" at:

800-671-FORE or 412-635-3700

3. You may FAX your questions to "support" at:

412-742-7900

4. You may send questions, via U.S. Mail, to:

FORE Systems, Inc. 1000 FORE Drive Warrendale, PA 15086-7502

Technical support for non-U.S.A. customers should be handled through your local distributor.

No matter which method is used for technical support, please be prepared to provide your support contract ID number, the serial number(s) of the product(s), and as much information as possible describing your problem/question.

Typographical Styles

Throughout this manual, all specific commands meant to be entered by the user appear on a separate line in bold typeface. In addition, use of the Enter or Return key is represented as <ENTER>. The following example demonstrates this convention:

cd /usr <ENTER>

File names that appear within the text of this manual are represented in the following style: "...the fore_install program installs this distribution."

Command names that appear within the text of this manual are represented in the following style: "...using the flush-cache command clears the bridge cache."

Subsystem names that appear within the text of this manual are represented in the following style: "...to access the bridge subsystem..."

Parameter names that appear within the text of this manual are represented in the following style: "...using <seg-list> allows you to specify the segments for which you want to display the specified bridge statistics."

Any messages that appear on the screen during software installation and network interface administration are shown in Courier font to distinguish them from the rest of the text as follows:

.... Are all four conditions true?

Important Information Indicators

To call your attention to safety and otherwise important information that must be reviewed to ensure correct and complete installation, as well as to avoid damage to the FORE Systems product or to your system, FORE Systems utilizes the following *WARNING/CAUTION/NOTE* indicators.

WARNING statements contain information that is critical to the safety of the operator and/or the system. Do not proceed beyond a **WARNING** statement until the indicated conditions are fully understood or met. This information could prevent serious injury to the operator, damage to the FORE Systems product, the system, or currently loaded software, and is indicated as follows:

WARNING!



Hazardous voltages are present. To reduce the risk of electrical shock and danger to personal health, follow the instructions carefully.

CAUTION statements contain information that is important for proper installation/operation. Compliance with **CAUTION** statements can prevent possible equipment damage and/or loss of data and are indicated as follows:

CAUTION



You risk damaging your equipment and/or software if you do not follow these instructions.

NOTE statements contain information that has been found important enough to be called to the special attention of the operator and is set off from the text as follows:



If you change the value of the LECS control parameters while the LECS process is running, the new values do not take effect until the LECS process is stopped, and then restarted.

Laser Notice

Class 1 Laser Product: This product conforms to applicable requirements of 21 CFR 1040 at the date of manufacture.

Class 1 lasers are defined as products which do not permit human access to laser radiation in excess of the accessible limits of Class 1 for applicable wavelengths and durations. These lasers are safe under reasonably foreseeable conditions of operation.

Every fiber optic interface on FORE Systems' network modules contains a Class 1 laser.



This Laser Notice section only applies to products or components containing Class 1 lasers.

Safety Precautions

For your protection, observe the following safety precautions when setting up equipment:

- Follow all warnings and instructions marked on the equipment.
- Ensure that the voltage and frequency of your power source matches the voltage and frequency inscribed on the equipment's electrical rating label.
- Never push objects of any kind through openings in the equipment. Dangerous
 voltages may be present. Conductive foreign objects could produce a short circuit
 that could cause fire, electric shock, or damage to your equipment.

Modifications to Equipment

Do not make mechanical or electrical modifications to the equipment. FORE Systems, Inc., is not responsible for regulatory compliance of a modified FORE product.

Placement of a FORE Systems Product

CAUTION



To ensure reliable operation of your FORE Systems product and to protect it from overheating, openings in the equipment must not be blocked or covered. A FORE Systems product should never be placed near a radiator or heat register.

Power Cord Connection

WARNING!



FORE Systems products are designed to work with single-phase power systems having a grounded neutral conductor. To reduce the risk of electrical shock, do not plug FORE Systems products into any other type of power system. Contact your facilities manager or a qualified electrician if you are not sure what type of power is supplied to your building.

WARNING!



Your FORE Systems product is shipped with a grounding type (3-wire) power cord. To reduce the risk of electric shock, always plug the cord into a grounded power outlet.

Preface

CHAPTER 1

Configuring PVCs

To establish a permanent communication link between two sites, it is necessary to establish permanent virtual circuits (PVCs) at each device along the communications path. This chapter discusses the creation of PVCs through the ATM Management Interface (AMI), a command-line user interface to the MSC-900.

1.1 General Concepts

Each ATM cell contains a virtual path identifier (VPI) and a virtual channel identifier (VCI) as part of its five-byte ATM header. The VPI and VCI are used to route the cell through the ATM network. When a switch fabric receives a cell, it examines the ATM header to determine the correct output port, VPI, and VCI for the cell. For example, an ATM switch fabric can be configured such that any cell received on port A1 with VPI |VCI| = 0 32 is switched to port B2 with VPI |VCI| = 0 35. The translation from input port, VPI, and VCI to output port, VPI, and VCI is achieved via a mapping table in the switch fabric's memory.

The VCI value of cells does not change as the cell is switched through the ATM network via a virtual path. In a single switch environment, a cell's VPI and VCI are translated only once, but in a multiple switch environment a cell's VPI and VCI are translated many times. It is important to remember that a cell's VPI and VCI are of local significance only (i.e., link-by-link). It is also important to note that virtual connections are unidirectional; that is, they are valid in one direction only. The VPI and VCI may change as the cell is switched through the network.

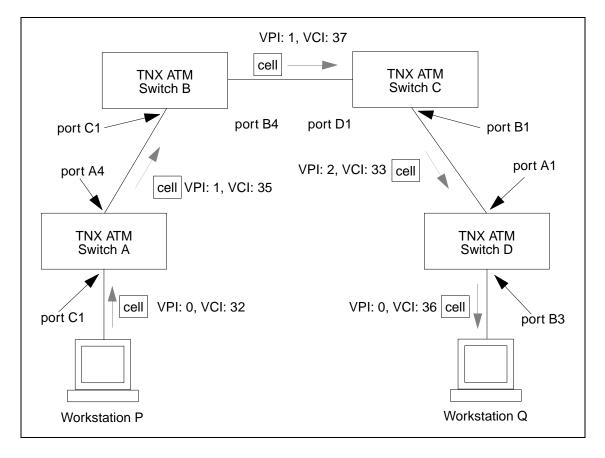


Figure 1.1 - The Cell

The mappings in an ATM network used to route cells from a source to a destination are generally referred to as virtual channels and virtual paths. The following sections explain how to create the necessary mappings to establish these virtual paths and virtual channels in an ATM network.

1.2 Virtual Paths

Virtual paths, which are carried within a physical transit medium (e.g., DS1, E1, DS3, E3, OC3c, or OC12c link), are used to establish connections between two nodes in an ATM network. Many virtual paths can be transmitted within a single physical link. Two types of virtual paths exist: virtual path connections (VPCs), also known as through paths, and originating/terminating paths, also known as virtual path terminators (VPTs). VPCs allow virtual paths to be cross-connected at a switch node while VPTs allow virtual channels (VCCs) to be cross-connected or switched at a switch node.

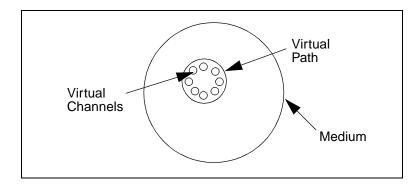


Figure 1.2 - Virtual Channels in a Virtual Path

A single virtual path can be used to route many virtual channels through the ATM network. Because a virtual path simply routes virtual channels through the network, a cell is guaranteed to have the same VCI when it exits the virtual path as it had when it entered the virtual path.

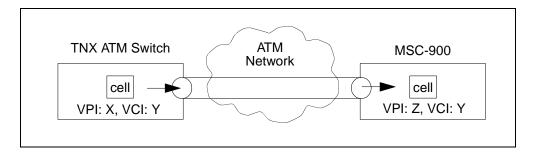


Figure 1.3 - An Example of a Virtual Path

The VCI value of cells does not change as the cell is switched through the ATM network via a through path. Each virtual path must originate at a switch fabric, pass through zero or more switch fabrics and terminate at another switch fabric. The origination and termination points are referred to as originating and terminating paths. Virtual paths are switched through switch fabrics via through paths. Virtual paths are made up of an originating path, zero or more through paths, and a terminating path.

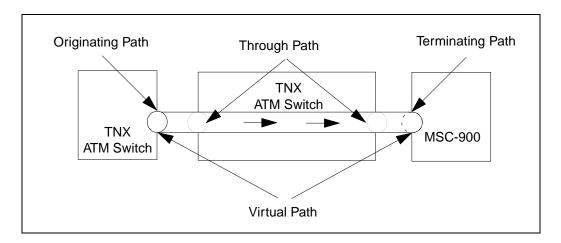


Figure 1.4 - Composition of a Virtual Path

1.2.1 Through Paths

Through paths route an entire VPC through an ATM switch fabric. When a cell is received by a switch fabric on a through path, the VPI is examined to determine the output port and VPI. The VCI component of the ATM header remains unchanged and can have any value. So, all of the channels within the through path are switched correctly without altering the VCI value of cells on these channels.

Four parameters are needed to define a through path on an MSC-900: input port, input VPI, output port, and output VPI. Through paths are represented as follows:

The VCI value remains unchanged when cells are switched via a through path. For example, the through path $A4 \mid 10 \rightarrow B4 \mid 20$ maps cells received on port A4 with VPI: 10 and any VCI to port B4 with VPI: 20 and the same VCI.

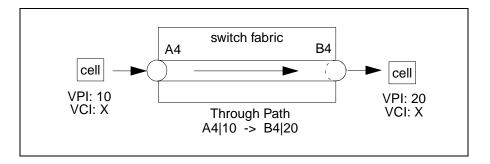


Figure 1.5 - An Example of a Through Path

By definition, through paths only switch cells in one direction; they are unidirectional. For example, switch fabric X is configured with the through path B1 | 20 -> C1 | 20. If a cell is received on port C1 with VPI: 20, it is not transmitted on port B1 with a new VPI: 20. In order for this to happen, the through path C1 | 20 -> B1 | 20 must exist as well. Since through paths are unidirectional, two through paths are necessary for bidirectional communication.

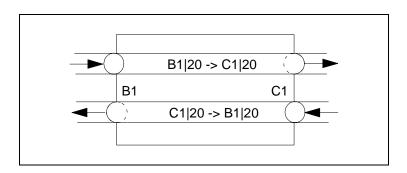


Figure 1.6 - Through Paths are Unidirectional

1.2.2 Originating and Terminating Paths

As previously noted, originating and terminating paths (also called virtual path terminators) are points at which a virtual path originates and terminates. For example, if a virtual path exists from switch fabric A to switch fabric B, then there must be an originating path on switch fabric A and a terminating path on switch fabric B.

An originating path is defined by two parameters: output VPI and output port. Similarly, a terminating path is defined by the parameters: input VPI and input port. Because originating and terminating paths do not define the way cells are switched through an ATM switch fabric, virtual channels must exist to switch cells from a terminating path to an originating path. (See the section about virtual channels for more information). Originating and terminating paths are the endpoints of virtual paths and are used primarily for bandwidth allocation.

The bandwidth allocated to originating and terminating paths is used to control the amount of virtual channel (VCC) bandwidth entering or leaving a virtual path. The total guaranteed bandwidth used by virtual channels on an originating path or a terminating path cannot exceed the amount of bandwidth allocated to that path. For example, as illustrated in Figure 1.7, if each of the four virtual channels shown is using 10 Mbps of bandwidth, then the originating and terminating paths must have at least 40 Mbps of bandwidth allocated.



UBR traffic bandwidth, which is a "best effort" service class, is not limited by the VP's allocated bandwidth since its bandwidth is not guaranteed. Actual UBR VCC traffictransmitted within a VP may exceed the VP's allocated bandwidth.

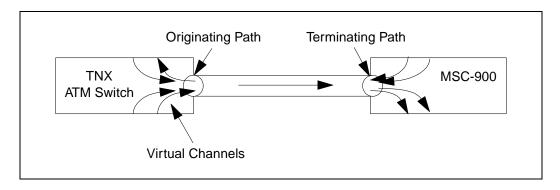


Figure 1.7 - Using Originating and Terminating Paths for Bandwidth Allocation

1.3 Listing Virtual Paths

1.3.1 Listing Through Paths

By logging in to AMI (see Chapter 1 of the *ForeThought-SP AMI Reference Manual* for information about logging into AMI), it is possible to display either all of the existing through paths on an individual switch fabric or all of the existing through paths on a specified port. To list all of the existing through paths on an individual switch fabric, enter the following parameters:

configuration vpc show

Input		Output						
	Port	VPI	Port	VPI	UPC	Prot	Name	
	3B1	40	3B4	40	0	pvc	customer_a	
	3B1	75	3B5	75	0	pvc	customer_b	
	3B2	95	3B3	95	0	pvc	customer_e	
	3B6	62	3B2	62	0	pvc	customer_c	
	3B6	68	3B3	68	0	pvc	customer_d	

The fields in this display are defined as follows:

Input Port Shows the number of the	the input port of the through
---	-------------------------------

path.

Input VPI Shows the input virtual path.

Output Port Shows the number of the output port of the through

path.

Output VPI Shows the output virtual path.

UPC Shows the integer index that refers to a specific traffic

contract assigned to this through path. See Section 2.21 in the *ForeThought-SP AMI Reference Manual*.

Prot Shows pvc for a permanent virtual circuit.

Name Shows the user-assigned name which helps to

identify this through path uniquely.

Configuring PVCs

To list advanced options about all of the existing virtual (through) paths, enter the following parameters:

configuration vpc show advanced

Input			Outpu	.t		
	Port	VPI	Port	VPI	Shape	ConType
	3B1	40	3B4	40		N/A
	3B1	75	3B5	75		N/A
	3B2	95	3B3	95		tran-tran-pmp
	3B6	62	3B2	62		tran-tran-pp
	3B6	68	3B3	68		N/A

The fields in this display are defined as follows:

Input Port Shows the number of the input port of the path.

Input VPI Indicates the input virtual path.

Output Port Shows the number of the output port of the path.

Output VPI Shows the output virtual path.

Shape Indicates whether or not traffic shaping has been

enabled for this path.

ConType Indicates the connection type for the endpoints of

this path with respect to a particular network. Orig (originating) means that the ingress/egress endpoint of the path is connected to the source node which is outside the network, tran (transit) means that the ingress/egress endpoint of the path is connected to a node within the network, and term (terminating) means that the ingress/egress endpoint of the path is connected to the destination node which is outside the network. pp means this is labelled as a point-to-point path, pmp means this is labelled as a multipoint-to-point path. mppm means this is labelled

as a multipoint-to-multipoint path.

1.3.2 Listing Originating and Terminating Paths

By logging in to AMI (see Chapter 1 of the *ForeThought-SP AMI Reference Manual* for information about logging into AMI), it is possible to display either all of the existing originating and terminating paths on an individual switch fabric or on a specified port. To list all of the existing originating and terminating paths on an individual switch fabric, enter the following parameters:

configuration vpt show

Input		Outpu	.t							
	Port	VPI	Port	VPI	ResBW	CurBW	MinVCI	MaxVCI	VCs	Protocol
	3B1	0	termi	nate	N/A	0.8K	1	511	6	pvc
	3B2	0	termi	nate	N/A	0.8K	1	511	6	pvc
	3B3	0	termi	nate	N/A	0.8K	1	511	6	pvc
	3B3	99	termi	nate	N/A	0.8K	1	511	2	pvc
	3B4	0	termi	nate	N/A	0.8K	1	511	6	pvc
	3B4	88	termi	nate	N/A	0.8K	1	511	2	pvc
	3B5	0	termi	nate	N/A	0.8K	1	511	6	pvc
	3B6	0	termi	nate	N/A	0.8K	1	511	6	pvc
	3E1	0	termi	nate	N/A	0.0K	0	7167	4	pvc
	3E2	0	termi	nate	N/A	0.0K	0	7167	4	pvc
	3E4	0	terminate		N/A	0.0K	0	7167	4	pvc
	3CTL	0	termi	nate	N/A	0.0K	1	1023	35	pvc
	origi	nate	3B1	0	N/A	0.0K	1	511	6	pvc
	origi	nate	3B2	0	N/A	0.0K	1	511	6	pvc
	origi	nate	3B3	0	N/A	0.0K	1	511	6	pvc
	origi	nate	3B3	99	N/A	0.0K	1	511	2	pvc
	origi	nate	3B4	0	N/A	0.0K	1	511	6	pvc
	origi	nate	3B4	88	N/A	0.0K	1	511	2	pvc
	origi	nate	3B5	0	N/A	0.0K	1	511	6	pvc

Press return for more, q to quit: q

The fields in this display are defined as follows:

Input Port Shows the number of the input port of the vpt.

Shows originate if it is an originating path.

Input VPI Shows the input virtual path number.

Output Port Shows the number of the output port of the vpt.

Shows terminate if it is a terminating path.

Output VPI Shows the output virtual path number.

ResBW	Shows the maximum amount of bandwidth (in				
	Kbps) that is reserved for the virtual channels using				
	this vpt. A value of N/A indicates that this path is an				
	elastic path. Elastic paths allocate and de-allocate				
	bandwidth for their channels from the link.				
CurBW	Shows the amount of bandwidth (in Kbps) that is				

CurBW Shows the amount of bandwidth (in Kbps) that is being used by the virtual channels using this vpt.

MinVCI Shows the bottom number for the range of VCIs that are reserved for VCCs on this virtual path terminator. The default is 1.

MaxVCI Shows the top number for the range of VCIs that are reserved for VCCs on this virtual path terminator. The default is 511.

VCs Shows the number of virtual channels that are currently using this vpt.

Protocol Shows pvc for a permanent virtual circuit.

To list all of the advanced options about all of the existing virtual path terminators, enter the following parameters:

configuration vpt show advanced

Input		Output				
Port	VPI	Port	VPI	Shape	VBROB	BuffOB
3B1	0	termi	nate	N/A	N/A	N/A
3B2	0	termi	nate	N/A	N/A	N/A
3B3	0	termi	nate	N/A	N/A	N/A
3B3	99	termi	nate	N/A	N/A	N/A
3B4	0	termi	nate	N/A	N/A	N/A
3B4	88	termi	nate	N/A	N/A	N/A
3B5	0	termi	nate	N/A	N/A	N/A
3B6	0	termi	nate	N/A	N/A	N/A
3E1	0	termi	nate	N/A	N/A	N/A
3E2	0	termi	nate	N/A	N/A	N/A
3E4	0	termi	nate	N/A	N/A	N/A
3CTL	0	termi	nate	N/A	N/A	N/A
originate		3B1	0		100	100
origi	nate	3B2	0		100	100
origi	nate	3B3	0		100	100
origi	nate	3B3	99		100	100

Press return for more, q to quit: q

The fields in this display are defined as follows:

Input Port Shows the number of the input port of the vpt.

Shows originate if it is an originating path.

Input VPI Indicates the input virtual path number.

Output Port Shows the number of the output port of the vpt.

Shows terminate if it is a terminating path.

Output VPI Shows the output virtual path number.

Shape Indicates whether or not traffic shaping has been

enabled for this originating vpt.

VBROB Shows the bandwidth overbooking level assigned to

this vpt, specified as a percentage. Valid values are integers from 1 to 99,999. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater

than 100 denote overbooking.

BuffOB Shows the buffer overbooking level assigned to this

vpt, specified as a percentage. Valid values are integers greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking.

Values greater than 100 denote overbooking.



For more information about setting overbooking parameters (VBROB and BuffOB), see Section 1.5.2 on page 2.21

1.5.2 on page 3-21.

1.4 Virtual Channels

Virtual channels "ride" inside of virtual paths. The combination of the two specifies a virtual connection. On a switch fabric, each virtual channel switches cells with a specific VPI and VCI received on a specific port to another port with a new VPI and a new VCI. Unlike through paths, which carry one or more VCCs, virtual channels describe a single virtual connection between two endpoints.

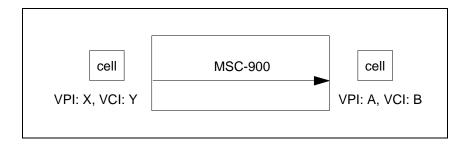


Figure 1.8 - An Example of a Virtual Channel

Six parameters are needed to define a virtual channel: input port, input VPI, input VCI, output port, output VPI, and output VCI. Virtual channels are represented by the following notation:

Virtual channels switch cells using both the VPI and VCI values. Both the VPI and VCI values may change when a cell is switched via a virtual channel. For example, the virtual channel $C2 \mid 1 \mid 20 \rightarrow D2 \mid 9 \mid 25$ switches cells received on port C2 with VPI: 1 and VCI: 20 such that they are transmitted out port D2 with VPI: 9 and VCI: 25.

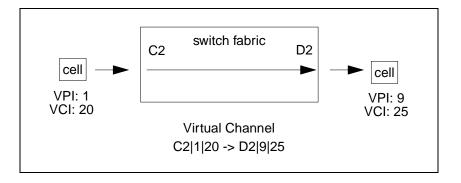


Figure 1.9 - Example of a Virtual Channel

In order to establish two-way communications between two ports on a switch fabric, two virtual channels are necessary because virtual channels are unidirectional. For example, switch fabric A is configured with the virtual channel C3 | 7 | 12 -> D1 | 8 | 2. If a cell is received on port D1 with VPI: 8 and VCI: 2, it is not transmitted out port C3 with VPI: 7 and VCI: 12. An additional channel, namely D1 | 8 | 2 -> C3 | 7 | 12, would have to exist.

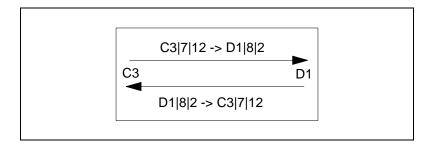


Figure 1.10 - Virtual Channels are Unidirectional

Before a virtual channel can be created, the corresponding terminating and originating paths must exist. For example, before the channels shown on the switch fabric in Figure 1.11 can be created, the terminating path $C3 \mid 3$ must exist.

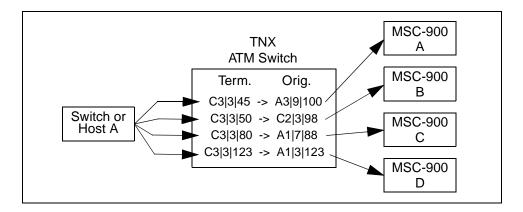


Figure 1.11 - Virtual Channels Created on Terminating Path C3 | 3

Similarly, before the virtual channels shown in Figure 1.12 can be created, the originating path $C2 \mid 2$ must exist.

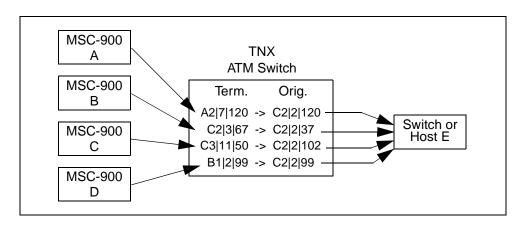


Figure 1.12 - Virtual Channels Created on Originating Path C2 | 2

Furthermore, in these examples, the terminating path $C3 \mid 3$ and originating path $C2 \mid 2$ must have enough bandwidth allocated to support the total bandwidth used by the virtual channels (see Figure 1.7).

1.4.1 Soft Permanent Virtual Circuits

Soft Permanent Virtual Circuits (SPVCs) are connections that go across multiple switch fabrics. An SPVC looks like a PVC at the local and remote endpoints with an SVC (Switched Virtual Circuit) in the middle. SVCs are channels established on demand by network signalling. Similar to a dialed telephone call, SVCs transport information between two locations and last only for the duration of the transfer.

SPVCs are more robust than PVCs. If a link carrying a PVC goes down, then the PVC goes down. If a link carrying a SPVC goes down and there is an alternate route, then the end switch fabrics of the SPVC automatically reroute the SPVC around the failed link.

In the diagram below, links exist between the TNX switches and the MSC-900. The endpoints exist at switch A and at MSC-900 A. If a link goes down between switch A and switch B, an SPVC can reroute the cell (via an SVC) through switch C.

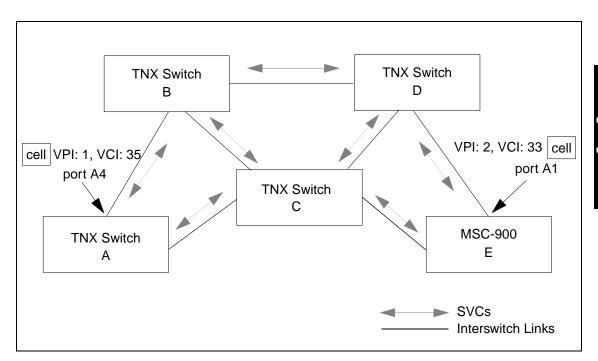


Figure 1.13 - The Path of a Cell Via SPVCs

Listing Virtual Channels 1.4.2

By logging in to AMI (see Chapter 1 of the ForeThought-SP AMI Reference Manual for information about logging into AMI), you can display either all of the existing virtual channels on an individual switch fabric or on a specified port. To list all of the virtual channels on an individual switch fabric, enter the following parameters:

configuration vcc show

Input		Ou	.tput					
Port	VPI	VCI	Port	VPI	VCI	UPC	Protocol	Name
3B1	0	5	3CTL	0	49	0	uni	N/A
3B1	0	14	3CTL	0	48	0	spans	N/A
3B1	0	15	3CTL	0	47		spans	N/A
3B1	0	16	3CTL	0	50		uni	N/A
3B1	0	100	3B4	0	100	0	pvc	N/A
3B2	0	5	3CTL	0	53	0	uni	N/A
3B2	0	14	3CTL	0	52	0	spans	N/A
3B2	0	15	3CTL	0	51		spans	N/A

Press return for more, q to quit: q

The fields in this display have the following meanings:

Input Port	Lists the incoming port number of the virtual channel.					
Input VPI	Shows the incoming virtual path number.					
Input VCI	Indicates the incoming virtual channel number.					
Output Port	Lists the outgoing port number of the virtual channel.					
Output VPI	Shows the outgoing virtual path number.					
Output VCI	Indicates the outgoing virtual channel number.					
UPC	Shows the integer index that refers to the specific UPC traffic contract assigned to this VCI.					
Protocol	Displays what type of protocol is running on this channel, which can be spans, pvc, uni or spvc.					
Name	Shows the user-assigned name which helps to identify this channel uniquely. If no name has been assigned, displays ${\tt N/A}$.					

1.5 Creating PVCs and SPVCs

FORE's ATM network modules provide ATM transmission connectivity, while its intelligent network modules, such as the Circuit Emulation Services (CES) network module, provide adaptation for ports carrying one transmission format (e.g., TDM) to ATM cells. This section describes how to create the following from one ATM port to another:

- · a permanent virtual path (through path)
- a permanent virtual path terminator (originating or terminating path)
- · a permanent virtual channel through the network
- a soft, permanent virtual channel through the network.

This section assumes that the physical port parameters of the switches have already been configured and that the ATM network module traffic models have been set appropriately. For more information about these configurations, see XXX.



When these paths and channels are created, a command is entered automatically into the current configuration database (CDB), meaning that these paths and channels are created every time the switch control processor (SCP) is restarted. The CDB can not be edited. Paths and channels can, however, be deleted manually via AMI. The CDB should be backed up frequently.

1.5.1 Creating a Through Path

To create a through path, log in to AMI (see Chapter 1 of the *ForeThought-SP AMI Reference Manual* for information about logging into AMI), and enter the following command:

```
conf vpc new <iport> <ivpi> <oport> <ovpi> [-upc <index>] [-name <name>]
```

The optional parameters for call records and VP shaping (using Series C network modules) are as follows:

Configuring PVCs

The parameters for creating through paths have the following meanings:

iport Indicates the incoming port number.

ivpi Indicates the incoming virtual path number.

oport Indicates the outgoing port number.

ovpi Indicates the outgoing virtual path number.

-upc <index> Indicates the integer index that refers to a specific

traffic contract. If no index is specified, then no traffic policing takes place on this VPI. It is assigned a UPC index of 0, and all traffic on this VPI is treated as UBR traffic. This is the default. (See Section 1.7.3 for more

information.)

-name <name> Indicates the name you want to assign to this

through path to help identify it uniquely. It is most useful for billing purposes so you can identify which paths are being used by which customers. Can be up

to 32 ASCII characters long.

NOTE

Terminating and originating paths can not be created across the intra-fabric ports on an MSC-900; only through paths can be created across the intra-fabric ports as shown in the following example.

The following is an example of how to create a virtual path which specifies a name:

```
localhost::configuration vpc> new 3b1 75 3b5 75 -name customer_b
```

The following is an example of how to create a virtual path on an MSC-900. To create a through path going in port 2A1, vpi 1 on on the switch board installed in slot 2 and going out port 4B1, vpi 1 on on the switch board installed in slot 4, enter the following:

```
localhost::configuration vpc> new 2a1 1 2e4 1
localhost::configuration vpc> new 2e4 1 2a1 1
localhost::configuration vpc> new 4b1 1 4e2 1
localhost::configuration vpc> new 4e2 1 4b1 1
```

In the first line in the first pair, notice that the output port is 2E4. This is the intra-fabric port. The 2 means the connection is coming out of the switch board in slot 2 through the intra-fabric port. The E represents the intra-fabric port. The 4 means the connection is destined for switch board in slot 4. 2E4 then becomes the input port in the second line.

In the first line in the second pair, notice that the output port is 4E2. This is the intra-fabric port. The 4 means the connection is coming out of the switch board in slot 4 through the intra-fabric port. The E represents the intra-fabric port. The 2 means the connection is destined for switch board in slot 2. 4E2 then becomes the input port in the second line.

At the same time, a command is entered automatically into the current configuration database, which means that this virtual path is created every time the SCP is restarted.

1.5.1.1 Through Paths with Call Records

The following is an example of how to create a virtual path which specifies a name and a connection type:

localhost::conf vpc> new 3b6 62 3b2 62 -name customer_c -inctype tran - outctype tran

The advanced options for call records to be logged per VC are as follows:

-inctype (orig | tran | term)

Indicates the path connection type for the incoming path. For billing purposes, it denotes on which switch this path is arriving. Orig (originating) means that the ingress endpoint of the path is connected to the source node which is outside the network, tran (transit) means that the ingress endpoint of the path is connected to a node within the network, and term (terminating) means that the ingress endpoint of the path is connected to the destination node which is outside the network.

-outctype (orig | tran | term)

Indicates the path connection type for the outgoing path. For billing purposes, it denotes on which switch this path is leaving. Orig (originating) means that the egress endpoint of the path is connected to the source node which is outside the network, tran (transit) means that the egress endpoint of the path is connected to a node within the network, and term (terminating) means that the egress endpoint of the path is connected to the destination node which is outside the network.

pmp Indicates this is a point-to-multipoint path.

mpp Indicates this is a multipoint-to-point path.

mpmp Indicates this is a multipoint-to-multipoint path.



By indicating pmp, mpp, or mpmp, you are only assigning a label for record keeping purposes. The switch does not necessarily create the type of path you have specified. If you assign a connection type, but do not assign a pmp, mpp, or mpmp label, the switch assigns a label of pp (point-to-point).

1.5.1.2 Through Paths with VP Shaping

The advanced options for VP shaping using Series C network modules are as follows:

-shapeivpi <vpi>

Indicates the incoming VPI for this through path. When the traffic shaping port is not the port connected to the WAN, a through path must be created from the WAN port to the traffic shaping port. Cells arrive from the network at the traffic shaping port with this value equal to the VPI of the terminating path at the traffic shaping port.



If you want to shape traffic on more than two ports on a given network module, it is recommended that you set the traffic memory model to model number 5 for that network module. Please see the sections, Displaying Traffic Models for a Network Module, and, Setting Traffic Models on a Network Module, found earlier in this chapter for information about how to do this.

1.5.2 **Creating an Originating or Terminating Path**

To create an originating or terminating path (virtual path terminator), log in to AMI (see Chapter 1 of the ForeThought-SP AMI Reference Manual for information about logging into AMI), and enter the following parameters:

```
conf vpt new <port> <vpi> (term | orig) [-reserved <Kbs>]
             [-minvci <vci>] [-maxvci <vci>]
```

The -reserved, -minvci, and -maxvci parameters are optional for creating originating paths. The advanced traffic management options for creating originating paths are as follows:

```
[-shapeovpi <vpi>] [-vbrob <percent>] [-vbrbuffob <percent>]
```

The parameters used for creating a virtual path terminator have the following meanings:

Indicates the port number for this vpt. port

Indicates the path number for this vpt. vpi

Specifies that the vpt to be created is a terminating term path. If this option is not used, an elastic path is

created. Elastic paths allocate and de-allocate bandwidth for their channels from the link.

Specifies that the vpt to be created is an originating orig

path.

The optional parameters used for creating originating paths have the following meanings:

Indicates the amount of bandwidth specified in Kbps -reserved <Kbs>

that you want to reserve on this originating vpt.

-minvci <minvci> Indicates the minimum VCI value for the range of

VCIs to be reserved for VCCs on this virtual path

terminator. The default is 1.

Indicates the maximum VCI value for the range of -maxvci <maxvci>

VCIs to be reserved for VCCs on this virtual path

terminator. The default is 1.

The advanced traffic management options for vpts have the following meanings:

-shapeovpi <vpi> Indicates the output port of a traffic shaping originating vpt. Setting this value configures traffic shaping on the

originating path. Cells bound for the network leave the traffic shaping port with this VPI. When the traffic shaping port is the WAN port, this value equals the input

VPI of the originating path. If the traffic shaping port is

not the WAN port, this value equals the input VPI of the through path from the shaping port to the WAN port.

-vbrob <percent>

Indicates the bandwidth overbooking level assigned to this vpt, specified as a percentage. Enter an integer value from 1 to 99,999. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.

-vbrbuffob <percent>

Indicates the buffer overbooking level assigned to this vpt, specified as a percentage. Enter an integer value greater than or equal to 1. The default is 100, which means that no overbooking has been defined. Values less than 100 cause underbooking. Values greater than 100 denote overbooking.



Bandwidth and/or buffer overbooking are used when the number of VBR VCCs configured across an originating or terminating VPC exceeds the guaranteed capacity of the VP. While overbooking removes service guarantees (in the case when all users transmit at full contract rate simultaneously), a network can be engineered in a more cost-effective manner if it is assumed that all VCC sources do not transmit simultaneously.

Guideline: An initial buffer overbooking value of 500-700% is recommended until live traffic patterns and VPC utilization can be measured. If bandwidth is still under-utilized, a higher VBROB setting can be used.

The following is an example of how to create a terminating path:

```
localhost::configuration vpt> new 3b3 99 term
Would you like to create the originating side also [y]? y
```

The following is an example of how to create a originating path:

```
localhost::configuration vpt> new 3b4 88 orig
Would you like to create the terminating side also [y]? y
```

1.5.2.1 Terminating a PVC at a Switch

Sometimes it is necessary to create a PVC between a host and a switch fabric that is at a remote location. In this case, the PVC should be created from the host to the control port (CTL) of the SCP and vice versa.

Some additional configuration is necessary for communication to be established between the host and the switch fabric. The SCP needs an entry in its ATM ARP cache in order to send cells destined for the host with the correct VPI and VCI and to pass received cells with a specific VPI and VCI to IP. This configuration can be done using AMI as shown in the following subsections. (See the atmarp (8c) man page for more information.)

1.5.2.2 Creating ATM ARP Entries

To create a FORE IP PVC ARP entry, log in to AMI (see Chapter 1 of the ForeThought-SP AMI Reference Manual for information about logging into AMI). Data on this PVC is encapsulated using null encapsulation (also known as VC-based multiplexing) as specified in RFC-1483. Enter the following parameters:

configuration atmarp newforeip <host> <vpi> <vci> (4|5) [<interface>]

Indicates the IP address of the remote host. host

igv Indicates the virtual path number. This must be 0.

Indicates the virtual channel number. vci

4|5 Designates the connection's ATM Adaptation Layer (AAL) type.

interface Indicates the FORE IP interface to be used. The

default is asx0.

Once the parameters are entered, the entry is created instantly by the SCP. At the same time, a command is entered automatically into the current CDB which creates this ATM ARP entry each time the SCP is restarted.

To create a new Classical IP PVC ARP entry, log in to AMI. All data is sent LLC/SNAP encapsulated. Enter the following parameters:

configuration atmarp newclassicalip <host> <vpi> <vci> [<interface>]

Indicates the host IP address of the remote IP host endstation.

Indicates the virtual path number of the Classical IP vpi PVC.

 ${f vci}$ Indicates the virtual channel number of the Classical

IP PVC.

 $\begin{tabular}{ll} \textbf{Indicates the name of the Classical IP interface to be} \\ \end{tabular}$

used for this connection. The default is qaa0.

Once the parameters are entered, the Classical IP PVC ARP entry is created instantly by the SCP. At the same time, a command is entered automatically into the current CDB which creates this ATM ARP entry each time that the SCP is restarted.

1.5.2.3 Listing ATM ARP Entries

To verify that the ARP entries exist correctly for the outgoing PVC connection from the SCP to the host, log in to AMI (see Chapter 1 of the *ForeThought-SP AMI Reference Manual* for information about logging into AMI). To display the ATM ARP cache, enter the following parameters:

configuration atmarp show

IPaddress	If	VPI	VCI	AAL	Type	Direction
198.29.22.9	asx0	0	63	aal5	foreIpSVC	pending
198.29.22.15	asx0	0	231	aal5	foreIpSVC	pending
198.29.22.37	asx0	0	65	aal34	foreIpSVC	pending
IPaddress	If	NSAP	Addres	ss		
198.29.17.3	qaa0	0x47.	0005.8	30.ffe10	0.0000.f21b.0138.	002048102754.00
198.29.17.10	qaa0	0x47.	0005.8	30.ffe10	0.0000.f21b.0137.	002048100be6.00
198.29.17.15	qaa0	0x47.	0005.8	30.ffe10	0.0000.f21b.0137.	00204810048d.00
198.29.17.52	qaa0	0x47.	0005.8	30.ffe10	0.0000.f21b.0138.	0020481b0138.00

The fields in this display have the following meanings:

IPaddress	Indicates the IP address for this connection.						
If	Shows the name of the IP interface for this connection.						
VPI	Displays the virtual path number.						
VCI	Displays the virtual channel number.						
AAL	Displays the AAL type of the given connection.						
Туре	Lists the kind of connection. Can be foreIpPVC, foreIpSVC, classicalIpPVC, or classicalIpSVC.						
Direction	Pending means that a connection has not (yet) been established. Incomplete means that the IP-to-ATM address mapping is not yet known for the given IP address.						
NSAP Address	Lists the NSAP address for this connection.						

1.5.3 **Creating a Virtual Channel**

To create a new virtual channel, log in to AMI and enter the following parameters:

```
conf vcc new <iport> <ivpi> <ivci> <oport> <ovpi> <ovci>
              [-upc <index>][-name <name>]
```

The advanced options for call records are as follows:

```
-inctype (orig | tran | term) -outctype (orig | tran | term)
                       [[qmqm|qqm|qmq]]
```

The parameters for this command have the following meanings:

iport

Indicates the incoming port number. Indicates the incoming virtual path number. ivpi Indicates the incoming virtual channel number. ivci oport Indicates the outgoing port number. ovpi Indicates the outgoing virtual path number. Indicates the outgoing virtual channel number. ovci Indicates the integer index that refers to a specific -upc<index> UPC traffic contract. If no index is specified, then no traffic policing takes place on this VCI. It is assigned a UPC index of 0, and all traffic on this VCI is treated as UBR traffic. This is the default. (See Section 1.7.3

for more information.)

Indicates the name you want to assign to this channel -name <name> to help identify it uniquely. It is most useful for

billing purposes so you can identify which channels are being used by which customers. Can be up to 32

ASCII characters long.

The advanced options for call records to be logged per VC are as follows:

-inctype (orig | tran | term)

Indicates the channel connection type for the incoming channel. For billing purposes, it denotes on which switch this channel is arriving. Orig (originating) means that the ingress endpoint of the channel is connected to the source node which is outside the network, tran (transit) means that the ingress endpoint of the channel is connected to a node within the network, and term (terminating) means that the ingress endpoint of the channel is connected to the destination node which is outside the network.

-outctype (orig | tran | term)

Indicates the channel connection type for the outgoing channel. For billing purposes, it denotes on which switch this channel is leaving. Orig (originating) means that the egress endpoint of the channel is connected to the source node which is outside the network, tran (transit) means that the egress endpoint of the channel is connected to a node within the network, and term (terminating) means that the egress endpoint of the channel is connected to the destination node which is outside the network.

pmp

Indicates this is a point-to-multipoint channel.

mpp

Indicates this is a multipoint-to-point channel.

mpmp

Indicates this is a multipoint-to-multipoint channel.



By indicating pmp, mpp, or mpmp, you are only assigning a label for record keeping purposes. The switch does not necessarily create the type of channel you have specified. If you assign a connection type, but do not assign a pmp, mpp, or mpmp label, the switch assigns a label of pp (point-to-point).

The following is an example of how to create a uni-directional virtual channel which specifies the call record connection type as "transit":

localhost::config vcc> new 3b1 0 100 3b4 0 100 -inctype tran -outctype tran

The following is an example of how to create a uni-directional virtual channel which has the name "customer_a" assigned to it:

```
localhost::configuration vcc> new 3b2 0 145 3b3 0 145 -name customer_a
```

The following is an example of how to create a bi-directional virtual channel on an MSC-900. To create a vcc going in port 2A1, vpi 0, vci 100 on the switch board installed in slot 2 and going out port 4B1, vpi 0, vci 100 on the switch board installed in slot 4, enter the following:

```
localhost::configuration vcc> new 2a1 0 100 2e4 0 100 localhost::configuration vcc> new 2e4 0 100 2a1 0 100 localhost::configuration vcc> new 4b1 0 100 4e2 0 100 localhost::configuration vcc> new 4e2 0 100 4b1 0 100
```

In the first line in the first pair, notice that the output port is 2E4. This is the intra-fabric port. The 2 means the connection is coming out of the switch board in slot 2 through the intra-fabric port. The E represents the intra-fabric port. The 4 means the connection is destined for switch board in slot 4. 2E4 then becomes the input port in the second line.

In the first line in the second pair, notice that the output port is 4E2. This is the intra-fabric port. The 4 means the connection is coming out of the switch board in slot 4 through the intra-fabric port. The E represents the intra-fabric port. The 2 means the connection is destined for switch board in slot 2. 4E2 then becomes the input port in the second line.

Once the parameters are entered, the virtual channel is created instantly by the SCP. At the same time, a command is entered automatically into the current configuration database, which means that this virtual channel is created every time the SCP is restarted.

1.5.4 Creating a SPANS SPVC

To create a SPANS SPVC, you must configure both ends of the connection concurrently on the two switch fabrics. This means you must have an AMI session open on both the local switch fabric and the destination switch fabric. To create a new SPANS SPVC, log in to AMI and enter the following parameters:

```
conf spvc spans new <port> <vpi> <vci> <dest-session> <dest-port> <dest-vpi> <dest-vci> \ [-peak <Kb/sec>] [(source | destination | bidirectional)]
```

The options for this command have the following meanings:

port Indicates the port number on the local switch fabric

on which this SPANS SPVC is to be created.

vpi Indicates the virtual path number on the local switch

fabric.

vci Indicates the virtual channel number on the local

switch fabric.

dest-session Indicates the name of the remote switch session.

dest-port Indicates the port number on the remote switch

fabric on which this SPANS SPVC is to be created.

dest-vpi Indicates the virtual path number on the remote

switch fabric.

dest-vci Indicates the virtual channel number on the remote

switch fabric.

-peak <Kb/sec> Indicates the amount of peak bandwidth allocated

for this SPANS SPVC, specified in kilobits per

second. The default value is 0.

source|destination|bidirectional Source means a unidirectional SPANS SPVC going

from the local switch fabric to the remote switch fabric will be created. Destination means a unidirectional SPANS SPVC going from the remote switch fabric to the local switch fabric will be created. Bidirectional means the pair of unidirectional SPANS SPVCs will be created. The default direction, if you

do not specify one, is bidirectional.

To cr

To create a bidirectional SPANS SPVC, you must either specify bidirectional, or you must set up two unidirectional SPANS SPVCs with one going

in each direction.

The following example depicts how to create a bidirectional SPVC from the local switch fabric (localhost) to a remote switch fabric (198.29.22.46). The user is logged into localhost.

```
localhost::> open 198.29.22.46 private

Opening a session for "198.29.22.46", please wait...
Connected to "198.29.22.46" (asx200bx).

198.29.22.46::> localhost

localhost::> configuration spvc spans new ?

usage: new <port> <vpi> <vci> <dest-session> <dest-port> <dest-vpi> <dest-vci> \[-peak <Kb/sec>] [(source | destination | bidirectional)]

localhost::configuration spvc spans> new 1c1 0 49 198.29.22.46 1b1 0 50
```

1.5.4.1 Displaying SPANS SPVC Information

This command allows you to display all of the SPANS SPVCs on an individual switch fabric. Enter the following parameters:

 localhost::configuration spvc spans> show

 Local
 Remote

 ID
 Port VPI VCI
 BW Direction
 ID
 Port VPI VCI Switch

 35664 1C1
 0 51
 0.0 bidirectional
 10427 1B1
 0 52 198.29.22.46

 65364 1C1
 0 49
 0.0 bidirectional
 42591 1B1
 0 50 198.29.22.46

The fields in this display are defined as follows:

y are defined as follows:				
Local ID	Indicates the unique number that the local switch fabric's SCP assigned to this SPANS SPVC when the it created.			
Local Port	Indicates the port number on the local switch fabric.			
Local VPI	Indicates the virtual path number on the local switch fabric.			
Local VCI	Indicates the virtual channel number on the local switch fabric.			
Local BW	Indicates the amount of peak bandwidth allocated for this SPANS SPVC, specified in kilobits per second.			
Remote ID	Indicates the unique number that the remote switch fabric's SCP assigned to this SPANS SPVC when it was created.			
Remote Port	Indicates the port number on the remote switch fabric.			
Remote VPI	Indicates the virtual path number on the remote switch fabric.			
Remote VCI	Indicates the virtual channel number on the remote switch fabric.			
Switch	Indicates the IP address or name of the remote switch fabric's SCP.			

The following is displayed if no SPANS SPVCS have been configured:

```
localhost::configuration spvc spans> show
No SPVC information is available
```

1.6 Traffic Types



The following sections contain informal definitions of the concepts presented. For a detailed understanding of these issues, please see the ATM Forum's UNI 3.1 and TM 4.0 Specifications.

Quality of Service (QOS) Management is based on the bandwidth parameters associated with a virtual connection and the class of service and ATM Adaptation Layer (AAL) used for that connection. In order to support voice, video, and data, the ATM Forum has defined four classes of service, or traffic types: Constant Bit Rate (CBR), Variable Bit Rate (VBR), Available Bit Rate (ABR), and Unspecified Bit Rate (UBR).

- At connection set-up time, traffic that uses a CBR parameter, such as a voice signal, makes a request for a dedicated Peak Cell Rate (PCR). Once the PCR is defined, the ATM network must be able to guarantee that amount of bandwidth for the duration of the connection.
- At connection set-up time, traffic that uses a VBR parameter, such as a video and data, makes a request for a dedicated PCR, Sustainable Cell Rate (SCR), and Maximum Burst Size (MBS). Once these cell rates are defined, the ATM network must be able to guarantee these rates for the duration of the connection.
- ABR traffic does not require timing relationships between source and destination.
 ABR provides no guarantees in terms of cell loss or delay, providing only "besteffort" service. ABR traffic sources adjust their transmission rate in response to
 information they receive describing the status of the network and its capability to
 successfully deliver data.
- UBR traffic, such as broadcast information and ARP messages, is also known as "best effort" service. UBR provides no bandwidth guarantees.

Because ATM is designed to provide a single network to transport these traffic classes, FORE's traffic policing and Connection Admission Control (CAC) schemes are vital to allowing this mix of traffic to flow smoothly.

1.7 Traffic Policing (Usage Parameter Control)

Traffic policing, also known as Usage Parameter Control (UPC), is a method ensuring fair allocation of network resources and of assessing the cells entering the switch for conformance with pre-established traffic bandwidth contracts. Those cells that exceed the specified contract are "tagged" or "dropped," depending on what is defined in the contract. This ensures that the connections with reserved bandwidth are not exceeding their reservations and using an unfair amount of network resources. FORE Systems' switches use a combination of "leaky bucket," or Generic Cell Rate Algorithm (GCRA) hardware in the switch fabric and user-configurable parameters in AMI to perform these policing functions.

1.7.1 Leaky Bucket Algorithm

The first important concept to understand is the leaky bucket algorithm. Leaky buckets are a mechanism by which cells entering the switch fabric are monitored for compliance with UPC traffic contracts that have been negotiated at connection set-up time. Before the leaky buckets are discussed, it is important to understand the parameters that are being measured by the buckets, as shown in Table 1.1. These parameters are informally defined as follows:

- Peak Cell Rate (PCR) the maximum number of cells per second.
- Cell Delay Variation Tolerance (CDVT) the tolerance for variation in the interarrival time of these cells, or the amount of jitter that can be accepted by the network.
- Sustainable Cell Rate (SCR) the average rate of cell transmission for this connection, taking bursting into account.
- Burst Tolerance (BT) the maximum amount of cells that can be transmitted

The leaky bucket algorithm is basically a timer which assesses if cells entering the switch fabric conform to the parameters listed above. As a cell arrives, the timer assesses if the cell is on time, late, or early. If the cell is determined to be on time or late (based on the traffic parameters), the cell is allowed to pass unchanged. If the cell is early (which, in turn, causes the cell stream to exceed the specified parameters), the cell is considered non-conforming and is either dropped or tagged (the CLP bit is set to 1), depending on the specified contract.

The first bucket in this analogy measures the PCR, or the rate at which the bucket drains. It also considers the CDVT, or the depth of the bucket. The second bucket measures the SCR, or the rate at which the bucket drains, and the BT, or the depth of the second bucket.

1.7.2 Non-conforming Cells: Tagging vs. Dropping

Second, it is important to understand the concept of tagging and dropping. Each ATM cell has a Cell Loss Priority (CLP) bit which indicates if the network can drop it under congested conditions. When the CLP bit is set to 0 (or CLP = 0), the cell is assessed for compliance with traffic parameters associated with the CLP = 0 stream. If the traffic parameters dictate that non-compliant cells should be "tagged," the CLP bit is set to 1 (or CLP = 1) by the UPC, which means that upon experiencing congestion further in the network, these CLP = 1 cells are dropped in preference to CLP = 0 cells.

Table 1.1 shows the various traffic contracts and user-configurable actions to be executed upon the detection of non-conforming cells.

Traffic Type	PCR	SCR	MBS	Bucket 1	Bucket 2	Policing on Switch Fabric	Action on non- conforming cells	
							Bucket 1	Bucket 2
CBR	PCR01			PCR01, CDVT		Yes	Drop CLP=0+1	
CBR0	PCR0, PCR01			PCR01, CDVT	PCR0, CDVT	Yes	Drop CLP=0+1	Drop CLP=0
VBR	PCR01	SCR01	MBS01	PCR01, CDVT	SCR01, BT01+CDVT	Yes	Drop CLP=0+1	Drop CLP=0+1
VBR0	PCR01	SCR0	MBS0	PCR01, CDVT	SCR0, BT0+CDVT	Yes	Drop CLP=0+1	Drop CLP=0

Table 1.1 - Summary of Traffic Contract Variables and Policing Actions

1.7.3 UNI 3.0 UPC Traffic Contract Parameters

The ATM Forum has defined different types of traffic contracts to be used in conjunction with these leaky buckets. The parameters that make up these types of contracts are defined as follows:

- pcr0 PCR for cells with CLP=0
- pcr01 PCR for the aggregate of the CLP=0 cells and the CLP=1 cells (all cells)
- scr0 SCR for cells with CLP=0
- scr01 SCR for the aggregate of the CLP=0 cells and the CLP=1 cells (all cells)
- mbs0 MBS for cells with CLP=0
- mbs01 MBS for the aggregate of the CLP=0 cells and the CLP=1 cells (all cells)
- tag sets CLP bit=1 for CLP=0 cells that fail the PCR0 test for CBR0 contracts or the SCR0/MBS0 test for VBR0 contracts

Configuring PVCs

The specific combinations of these parameters that make up the ATM Forum contracts are defined as follows:

- 1. cbr <pcr01>
- 2. cbr0 < pcr0 > < pcr01 > [tag]
- 3. vbr <pcr01> <scr01> <mbs01>
- 4. vbr0 < pcr01 > (scr0) < mbs0 > [tag]

The cbr cpcr01> contract is for CBR traffic. It only uses the first leaky bucket to assess the conformance to PCR of the aggregate of the CLP=0 cells and the CLP=1 cells. Cells which fail the PCR CLP=0+1 test are discarded.

The cbr0 <pcr0> <pcr0> <pcr01> [tag] contract is for CBR traffic. It uses the first leaky bucket to assess the conformance to PCR of the CLP=0 cells. It uses the second leaky bucket to assess the conformance to PCR of the aggregate of the CLP=0 and the CLP=1 cells. If the tag option is set, the cells which fail the PCR CLP=0 test are tagged as CLP=1 and passed on to the second leaky bucket to be tested for PCR CLP=0+1 conformance. Cells which fail the PCR test on CLP=0+1 are discarded. If the tag option is not set, cells which fail the PCR CLP=0 test are discarded and cells which fail the PCR test on CLP=0+1 are discarded.

1.7.4 AMI UPC Commands

AMI allows you to create a UPC contract using these combinations of traffic parameters. To create a UPC contract in AMI, enter the following parameters:

```
localhost::conf upc> new <index> ubr [aal5 [noPktDisc]] [ubrTagging] [-name <name>]
   OR
new <index> <UPC> [-cdvt <us>] [noGCRA] [aal5 [noPktDisc] [PPPol]] [-name <name>]
```

For UBR contracts, the parameters are defined as follows:

index

Indicates an arbitrary, user-assigned integer that represents this particular contract. This number is used as the <code>-upc</code> <code>-index></code> parameter when creating a VPC, VCC, or SPANS signalling path that you want this contract to assess.



When a UPC contract is created, it is not actually enabled until it is assigned to a VPC, VCC, or SPANS path. UPC indices are not assigned to VPTs.

ubr Indicates that this is a UBR connection.

The optional UBR parameters are as follows:

aal5 Indicates that the connection is using the AAL5

Adaptation Layer.

noPktDisc This optional parameter can only be used if the

connection is AAL5 (i.e., aal5 must be entered). The noPktDisc parameter suppresses the Early Packet Discard/Partial Packet Discard (EPD/PPD) operations (AAL5 packet discard) on the connection.

By default, EPD/PPD is enabled.

ubrTagging Entering ubrTagging means that all UBR traffic is

tagged (set to CLP=1) on this connection. If ubrTagging is not entered, then UBR traffic is not tagged on this connection. This command only applies to UBR traffic. By default, UBR traffic is not

tagged.

name Indicates an optional UPC name that provides a

simple, user-defined name for the UPC contract.

For CBR, VBR, or UBR contracts, the parameters are defined as follows:

index

Indicates an arbitrary, user-assigned integer that represents this particular contract. This number is used as the -upc <index> parameter when creating a VPC, VCC, or SPANS signalling path. UPC indices are not use on VPTs.



When a UPC contract is created, it is not actually enabled until it is assigned to a VPC, VCC, or SPANS path.

UPC

Indicates one of the four traffic contract types and their associated parameters, as shown below:

```
cbr <pcr01>
cbr0 <pcr0> <pcr01> [tag]
vbr <pcr01> <scr01> <mbs01>
vbr0 <pcr01> <scr0> <mbs0> [tag]
```

The optional parameters are as follows:



PVCs that use UPC contracts which contain [noGCRA], [aal5 [noPktDisc] [PPPol]], and [ubrTagging] options are valid only when the conf port gcrapolicing, conf port aal5packetdiscard, conf port pppolicing, and conf port ubrtagging commands are set to svcOn or svcOff. Use conf port show tm to check these settings.

-cdvt <us>

Indicates the Cell Delay Variation Tolerance (CDVT) associated with the peak cell rates, specified in microseconds. If the CDVT is not specified here, the default CDVT value associated with the port will be used. (See conf port show and conf port cdvt for more information).

noGCRA

Entering nogcra means that GCRA policing is disabled on CBR or VBR connections (depending on what is configured) using this contract. By default, GCRA policing is enabled.

aal5 Indicates that the connection (typically VBR) is using the AAL5 Adaptation Layer format.

noPktDisc This optional parameter can only be used if the connection is AAL5 (i.e., aal5 must be entered). The noPktDisc parameter suppresses the Early Packet Discard/Partial Packet Discard (EPD/PPD) operations (AAL5 packet discard) on the connection.

By default, EPD/PPD is enabled.

This optional parameter can only be used if the connection is AAL5 (i.e., aal5 must be entered). This parameter indicates that Partial Packet Policing is going to be performed on this connection. By default, Partial Packet Policing is disabled.

HDCOMPversion must be 1 or greater, otherwise AAL5 partial packet policing is not supported. To display this version, use the conf board show advanced command.

Indicates an optional UPC name that provides a simple, user-defined name for the UPC contract.

The following is an example of how to create a UPC contract:

name

PPPoI

localhost::conf upc> new 5 vbr0 500 200 250 -cdvt 1000 aal5 PPPol -name vbr0 upc

This example specifies a contract named "vbr0_upc", which is a VBR0 contract with an index of 5, a pcr01 of 500 cells/sec, an scr0 of 200 cells/sec, an mbs0 of 250 cells (or kilobits), a CDVT of 1,000 microseconds, and partial packet policing enabled.



For more information regarding traffic contracts, please refer to Table 5-7 in the ATM Forum UNI 3.0 Specification.

Configuring PVCs

CHAPTER 2

Configuring a Classical IP ATM Network

2.1 Introduction

This chapter describes how to design, configure, and maintain a Classical IP ATM network. The term classical indicates that the ATM network has the same properties as existing legacy LANs. That is, even though ATM technology allows for large, globally connected networks, for example, it is only used in the LAN environment as a direct replacement of existing LAN technology. The classical model of LANs connected through IP routers is maintained in ATM networks. RFC-1577 provides the standard for Classical IP over ATM.

Classical IP over ATM is different than IP in legacy LANs in that ATM provides a virtual connection environment through the use of Permanent Virtual Circuits (PVCs) and/or Switched Virtual Circuits (SVCs). SVC management is performed via the ATM Forum UNI 3.0 Specification, which specifies Q.2931. Q.2931 is a broadband signalling protocol designed to establish connections dynamically at the User-Network Interface (UNI). Q.2931 uses Service Specific Connection Oriented Protocol (SSCOP) as a reliable transport protocol, and all signalling occurs over VPI: 0, VCI: 5. Q.2931 connections are bidirectional, with the same VPI/VCI pair used to transmit and receive.

Once a Classical IP connection has been established, IP datagrams are encapsulated using IEEE 802.2 LLC/SNAP and are segmented into ATM cells using ATM Adaptation Layer type 5 (AAL5). Additionally, the default Maximum Transmission Unit (MTU) is 9,180 bytes (the SNAP header adds 8 more bytes) with a maximum packet size of 65,535 bytes. There is currently no support for IP broadcast datagrams or IP multicast datagrams in a Classical IP environment.

2.1.1 Logical IP Subnets

An important concept in Classical IP networks is that of a Logical IP Subnet (LIS). An LIS is a group of hosts configured as members of the same IP subnet (that is, they have the same IP network and subnetwork numbers). In this sense, one LIS can be equated to one legacy LAN. It is possible to maintain several overlaid LISs on the same physical ATM network. Therefore, in a Classical IP ATM network, placing a host on a specific subnet is a logical choice rather than a physical one. In this type of environment, communication between hosts in different LISs is only permitted by communicating through an IP router which is a member of both LISs (as per RFC-1577).

The number of LISs, and the division of hosts into each LIS, is purely an administrative issue. Limitations of IP addressing, IP packet filtering, and administrative boundaries may guide a manager into establishing several LISs onto a single ATM network. Keep in mind, though, that communication between LISs must occur through IP routers.

2.1.2 Classical IP Interfaces

In order to support routing between multiple LISs, the host adapter software allows a host to be configured as a member of (and a router between) up to four distinct LISs per physical ATM interface. Each LIS membership is through a separate Classical IP network interface. Existing system level IP routing configuration tools are used to control routing through each of the Classical IP interfaces in the same manner as routing among several physical interfaces. Note that even though each Classical IP interface associated with a given physical interface uses the same physical hardware, they are each configured separately with their own MTU, IP address, and ATM address.

The name of each of the Classical IP interfaces begins with qa. All of the Classical IP interfaces associated with physical unit zero have a as the next letter. All of the Classical IP interfaces associated with physical unit one have b as the next letter, and so forth. Finally, each Classical IP interface has its interface number as a suffix. As an example of the above naming convention, the name of the third Classical IP interface (unit 2) on physical unit one is qab2.

2.1.3 SPANS Interface

While each of the qa interfaces for a given physical interface is designed to support Classical IP using Q.2931 signalling, a SPANS interface (usually called fa, but user configurable) also exists for each physical interface. The fa interface supports FORE IP on top of SPANS signalling. FORE IP allows communication using AAL4 or AAL5 with no encapsulation, uses a broadcast ARP for SPANS address resolution, and supports direct communication of all hosts on a physical ATM network without the use of IP routers. Since SPANS and Q.2931 signalling use different VCIs, a host can simultaneously support FORE IP over SPANS as well as Classical IP over Q.2931 on the same physical interface.

As a result of standard IP routing, all traffic sent out an fa interface uses FORE IP, while all traffic sent out a qa interface uses Classical IP. Each of the fa interfaces should be assigned an IP address on a subnet different than the subnets of any of the qa interfaces. It is permissible to place multiple fa interfaces on the same subnet, and the driver load balances connections across these interfaces.

It is only necessary to configure the fa and qa interfaces if the specific service provided by that interface is required. A host sending only Classical IP would not need to configure the fa interfaces. Likewise, a host sending only FORE IP would not need to configure the qa interfaces. Both the fa and qa interfaces may be configured simultaneously, but they must be in separate subnets. Remember that Classical IP specific configuration changes can only be done with the qa devices, while SPANS specific configuration changes can only be done with the fa devices.

2.2 Address Registration and ILMI

Before a host can establish connections over a physical interface, the host must know the NSAP address for that interface. The primary purpose of Interim Local Management Interface (ILMI) is to discover and register these NSAP addresses dynamically.

2.2.1 NSAP Addresses

For private ATM networks, addresses uniquely identify ATM endpoints. The UNI 3.0 address format is modeled after that of an OSI Network Service Access Point, hence the name NSAP address.

Three address formats have been specified: DCC, ICD, and E.164. FORE implements the ICD ATM format. Per the UNI 3.0 specification, all private networks should accept initial call set-up messages containing ATM addresses with any of the approved formats and forward the calls as necessary.

An NSAP address consists of the following:

- a 13-byte network-side prefix The prefix is the NSAP prefix of the SCP to which the host is attached.
- a seven-byte user-side part This consists of the following:
 - a six-byte End System Identifier (ESI) The ESI is the unique IEEE MAC address of the interface.
 - a one-byte selector Although each Classical IP interface for a given physical interface uses the same prefix and ESI, the selector field is the part that indicates the number of the specific Classical IP interface. The selector field is 00 for qaa0, 01 for qaa1, 02 for qaa2, and 03 for qaa3.

2.2.2 Operating with ILMI Support

FORE Systems devices running *ForeThought-SP* software version 1.0 or greater provide support for ILMI. If ILMI is supported on all of the devices in a given network, when a switch control processor (SCP) boots up, ILMI enables the SCP to discover all of the hosts attached to it and to send its NSAP prefix associated with the port to those hosts dynamically. In return, the host prepends that prefix to its ESI and selector fields, forming a complete NSAP address. The host then notifies the SCP of its complete NSAP address. These registration SNMP messages are sent and received over AAL5 using VPI: 0, VCI: 16. Once ILMI registration has been completed, then connection setup can occur.

If a host changes network ports after an NSAP address has been registered for its interface, all existing connections are closed. If the new port is on a different SCP, a new NSAP address (with a different network address prefix) is registered. The host can then begin to establish new connections.

2.2.3 Operating without ILMI Support

If ILMI is not supported on a particular SCP or host in a given network, then the NSAP addresses must be manually configured. If another vendor's switch does not support ILMI, it can not supply an NSAP prefix to the hosts. Therefore, the user must assign a unique, valid prefix to the switch. Additionally, the same prefix should be used for all hosts in the LIS.

On the host, atmarp -n (8C) is used to configure the NSAP address for a specific interface. The SCP directly attached to this interface is then informed of this NSAP address/port combination through commands in the ATM Management Interface (AMI). Once the host and network have both been informed of this NSAP address/port pair, the host may begin signalling.

Once an interface's NSAP address has been set, it is possible to change it by first running atmarp -n to set the new address, and then bringing the interface down and back up (ifconfig gaxx down; ifconfig gaxx up) for the address change to take effect.

2.2.4 Configuration

The choice to use ILMI for address registration is made at software installation time. Since ILMI uses SNMP as its management protocol, the use of ILMI is tied into snmpd. The choice can be made to run FORE's SNMP agent and use ILMI (snmpd), run FORE's SNMP agent without using ILMI (snmpd -n), or just use ILMI (snmpd -i or ilmid -i). The commands to run the chosen daemon, or to configure a host's NSAP address statically when not running ILMI, are placed in the firmware download script.

2.3 ARP and ARP Servers

2.3.1 Theory

In order for a host to establish a connection to another host, it must first determine the other host's NSAP address. ATM ARP (ATM address resolution protocol) is the procedure used to resolve an IP address into an ATM address. Since the ATM standards do not currently support broadcast on an ATM LAN, address resolution is performed by direct communication with a special ARP server, rather than broadcasting ARP requests as is done in legacy LANs. Each LIS must have only one ARP server configured, but a single ARP server can be the server for several LISs.

Each host in an LIS must be configured with the NSAP address of the host providing ARP service for its LIS. On a host ARP server, the NSAP address of the ARP server can be obtained by running atmarp -z (remember to use the interface associated with the given LIS). The ARP server address is normally configured into each host at installation time (the appropriate commands are placed in the firmware download script), but it may be done at any time by running atmarp -p. If the ARP server address is configured in the firmware download script, it is persistent across reboots, but if it is configured using atmarp -p, it is not persistent across reboots.

Since only one ARP server can be functioning at a time in a given LIS, and since the ARP server's address is manually configured into each host, it is not possible to use multiple, redundant ARP servers to improve robustness. If an ARP server becomes nonfunctional, a new ARP server must be configured, and then each host within the LIS must be configured to use the new ARP server. To configure a new ARP server address on a host, run atmarp -p to set the new address, then bring the appropriate interface down and back up (ifconfig qaxx down; ifconfig qaxx up) for the change to take effect.

2.3.2 Configuring a Host to be an ARP Server

The following procedures list the required steps to configure a host as an ARP server. Before proceeding, several preconditions MUST be met. You must also obtain both the interface name and ARP server NSAP address.

The preconditions are as follows:

- 1. A FORE adapter must be installed in the network.
- 2. The FORE Systems software has been installed and the system has been successfully booted.
- 3. The network must contain devices that recognize ILMI and RFC-1577, such as FORE switches and adapters.
- 4. When the host was initially configured, Classical IP should have been specified. By specifying Classical IP, ILMI and RFC-1577 are automatically active in the *ForeThought* SP 1.0 and greater adapter code.

ARP server/client configurations are run on the Classical IP interfaces. Classical IP interface names are of the form qaan, where N is from 0 to 3. Once the prerequisites have been met, the steps for configuring a host as an ARP server are as follows:

1. If not already done during installation, configure the IP, netmask, and broadcast address for the interface using the ifconfig command as follows:

2. Using the host interface name from the previous step, determine the NSAP address of the interface name you wish to make the ARP server with the following command (assuming the interface name is qaa0):

/usr/etc/fore/atmarp -z gaa0

A typical response is shown below:

NSAP addr for gaa0 is 0x47.0005.80.ffe100.0000.f215.0ce5.002048100c47.03



If the response is all zeros, this indicates that the NSAP address has not been assigned. Check the other <code>qaaN</code> interfaces to see if any have been configured. If not, then manually assign an NSAP address as in step 1 above.

If you are using a windowed graphical user interface, keep the window with the NSAP address open because you need to cut and paste this address into several locations. If not, carefully write down the address so you can add it correctly to the other locations. If you are not registered with ILMI, you need to assign or manually configure an NSAP address to the interface.

With the preconditions met and the required information obtained, ARP servers may be created with the following steps. Two methods are shown. An ARP server can be configured for a one-time use, which is deactivated if the ARP server is rebooted. This method is shown first. An ARP server can also be permanently configured to be active on boot. This method is shown second.

2.3.2.1 Configuring a Host as an ARP Server for One-Time Use

- 1. On the host you wish to configure as an ARP server, log in as root.
- On the host, change to the directory in which the ATM boot file is stored. To determine the location of this directory, use your system specific find command and search for the location of the atmarp utility. The ATM boot files are stored in this directory. (See the list below for the names of the system specific boot files.)
- 3. From the command prompt in this directory, issue the following command:

```
atmarp -p <NSAP address> <interface name>
```

The location of the ATM boot file is a function of the type of host in which the FORE adapter has been installed. On your system, the files may have been copied to another location during the installation process. The following list shows the default locations:

Solaris

```
/etc/rcS.d/S99rc.sba200
```

SunOS

/usr/etc/fore/rc.sba200

- IRIX

/etc/init.d/fore atm

- HP-UX

/etc/fore atm

4. On each host requiring service from the ARP server, use the following command:

atmarp -p <ARP Server NSAP address> <interface name>

2.3.2.2 Configuring a Host as a Permanent ARP Server

To configure an ARP server so that it persists across reboots, use the following procedure:

- 1. On the host you wish to configure as an ARP server, log in as root.
- 2. In the directory in which the ATM boot file is stored, edit the ATM boot file. Note where this file is in the directory structure from the list of operating systems above. For example, for an ESA-200 adapter, the file name is fore_atm. In this file, add the following line:

```
atmarp -p <NSAP_address> <interface_name>
```

where the correct NSAP address and interface names are used.

3. On each host requiring service from the ARP server, in the directory in which the ATM boot file is stored, edit the boot file. For example, for an ESA-200 adapter, the file name is fore_atm. In this file, add the following line:

```
atmarp -p <ARP_server_NSAP_address> <interface_name>
```

where the correct NSAP address and interface names are used. For example:

```
atmarp -p 0x47.0005.80.ffe100.0000.f215.0ce5.002048100c47.03 qaa0
```

Once these changes have been made, the ARP server is active upon booting. Similarly, each host attached to the ARP server can access the ARP server upon booting.

2.3.2.3 Configuring a Hewlett-Packard Computer as an ARP Server

On a Hewlett-Packard host, use the program /usr/etc/fore/config_atm to configure the host as either being connected to the ARP server or being the ARP server itself. Prior to running config_atm, see Section 2.3.2 for instructions on how to determine the host NSAP address and interface names to be used. The discovery process is the same for all computers.

2.3.3 Configuring an MSC-900 to be an ARP Server

The MSC-900 has the capability of being an ARP server. To configure an MSC-900 as an ARP server, perform the following steps on <u>only one</u> of the SCPs:

1. On one of the SCPs, determine the NSAP address of that SCP for the relevant interface (qaa0 -> qaa3) using the following AMI command:

configuration atmarp getnsap <interface>

For example:

configuration atmarp getnsap qaa0

gaa0 NSAP address: 47000580ffe1000000f12400de0020481900de00

2. Set the NSAP address of the ARP server to be the NSAP address of the interface that you displayed in step 1 using the following AMI command:

conf atmarp arpserver set <NSAPaddress> [<interface>]

For example:

conf atmarp arpserver set 0x47000580ffe1000000f12400de0020481900de00 qaa0

2.3.4 Classical IP Operation

Once a host knows its own ATM address and the ATM address of its ARP server it attempts to establish a connection to the ARP server, which is used to send ARP requests and receive ARP replies. When the connection to the ARP server has been established, the ARP server sends an inverse ARP (InARP) request on the new VC to learn the host's IP address. When an InARP reply is received, the ARP server places that host's IP address to ATM address mapping in its ARP cache. Therefore, over time, the ARP server dynamically learns the IP-to-ATM address mappings of all the hosts in its LIS. It can then respond to ARP requests directed toward it for hosts in its LIS.



In order for a host to communicate with an ARP server, it must have learned its own ATM address and have been configured with the ATM address of the ARP server.

A host can not resolve the ATM addresses of hosts in its LIS unless it can communicate with its ARP server.

Since there is no mechanism for ARP servers to exchange mapping information with each other, it is imperative that each LIS be configured with only one ARP server.

When a host wants to communicate with another host in its LIS, it first sends an ARP request to the ARP server containing the IP address to be resolved. When an ARP reply is received from the ARP server, the host creates an entry in its ARP cache for the given IP address and stores the IP-to-ATM address mapping. This ARP cache entry is marked as complete. To ensure that all of the IP-to-ATM address mappings known by a certain host are up-to-date, hosts are required to age their ARP entries. A host must validate its ARP entries every 15 minutes (20 minutes on an ARP server).

A host validates its SVCs by sending an ARP request to the ARP server. A host validates its PVCs, and an ARP server validates its SVCs, by sending an InARP request on the VC. If a reply is not received, the ARP entry is marked invalid. Once an ARP entry is marked invalid, an attempt is made to revalidate it before transmitting. Transmission proceeds only when validation is successful. If a VC associated with an invalid ARP entry is closed, the entry is removed.

2.3.5 Operational Issues

Certain hosts in an LIS may not support Classical IP. It is still possible to communicate with these hosts (and for these hosts to communicate with one another) by using static ARP entries. If a host does not support Classical IP, its IP-to-ATM address mapping should be placed in its ARP server's cache as a static entry. This allows other hosts that do support Classical IP to contact their ARP server as usual and obtain the correct address mapping. If a host that does not support Classical IP wants to initiate connections, the IP-to-ATM address mappings of the destination hosts should be put in its ARP cache, again as static entries. By using static ARP entries in the above fashion, the ability for all hosts to communicate is maintained.

There are some restrictions on the number of hosts that can be maintained as static ARP entries. They are as follows:

- In the default configuration, a host can only have approximately 250 virtual connections open simultaneously. This means that an ARP server can only serve 250 clients, since each client must maintain a connection with its ARP server. This may be a limitation if the ARP server is servicing multiple LISs.
- On the host, it is possible to increase the number of connections allowed by first using cpath (8c) to delete path 0 for the given port, and then using cpath to recreate path 0 for the given port, but using the -v option to cpath to specify the number of connections to allow over that path.
- Hosts support a maximum of 1,024 connections.

2.4 Classical IP PVCs

2.4.1 Theory and Configuration

Normally, ATM connections in a Classical IP environment are established dynamically using Q.2931. ARP, ILMI, and Q.2931 all work together as described previously to set up an SVC. If a host from another vendor does not support Classical ARP or ILMI, it is still possible to set up an SVC using work-arounds. If a host or a switch in an LIS does not support Q.2931, however, it is not possible to establish an SVC. In this case, a Classical IP PVC can be used for communication.

On each of the hosts, atmarp -c is used to establish the PVC. An unused VPI/VCI pair must be chosen for each host. PVCs using the chosen VPI/VCI pairs must also be set up from each of the hosts to their connecting switch, and then on all of the switches between the two connecting switches.



Both the incoming and outgoing connections are set up simultaneously on the host, but they must be set up individually on the switches. The same VPI/VCI pair is used by a host to send on the PVC as well as receive on the PVC. The IP datagrams are sent over the PVC using AAL5 with LLC/SNAP encapsulation.

2.4.2 Revalidation and Removal

Normally, the device driver periodically checks that its PVCs are still established and functioning. A host revalidates a PVC every 15 minutes by sending InARP requests over the PVC, if the user specifies that revalidation should occur by choosing 1 as the revalidate option to atmarp -c. If the equipment attached to the FORE equipment supports revalidation, the user must choose 1 as the revalidate option to atmarp -c. If an InARP reply is not received, the revalidation fails, the PVC is marked invalid (as shown through atmarp -a), and communication over the PVC is no longer possible.

Once a PVC is marked invalid, an attempt is made to validate the PVC before transmitting. Transmission proceeds only when validation is successful. It is possible to disable this revalidation feature by specifying 0 as the revalidate option to atmarp -c. This is often desirable when the remote end of the PVC (such as a video camera) does not support Inarp.

A Classical IP PVC is removed on the host side using atmarp -r. Both the incoming and outgoing connections are removed simultaneously. The PVC must then be removed from each of the network switches involved.

2.5 Debugging

The atmarp utility on the host provides a number of useful options to aid communication debugging. Some of this information can also be displayed for the SCP using various AMI commands. If an equivalent AMI command is available, it is listed here.

- On the host, atmarp shows the IP address of the entry, the associated ATM address if the mapping is known, the VPI/VCI pair for an established connection, and various flags. The "Classical IP" flag indicates connections using Classical IP over Q.2931. Two other flags are especially important for connection state: "incomplete" indicates that the IP-to-ATM address mapping is not yet known for the given IP address; "pending" indicates that a connection has not (yet) been established. For the MSC-900, AMI command configuration atmarp show provides the same information.
- On the host, atmarp -z displays the ATM address of the given interface. If the host is using ILMI for address registration, the ATM address is displayed as nonzero when the host has successfully registered its address with the network. For the MSC-900, AMI command configuration atmarp getnsap interface provides the same information.
- On the host, atmarp -g displays the ATM address of the current ARP server for the LIS associated with the given interface. For the MSC-900, AMI command configuration atmarp arpserver show interface provides the same information.
- On the host, atmarp -t indicates whether this host is the ARP server for the LIS associated with the given Classical IP interface. For the MSC-900, AMI command configuration atmarp arpserver show interface provides the same information.
- On the host, atmarp -v displays the status of the ARP server connection. One of the first things to check in a trouble state is that a connection to the ARP server is fully established. If the ARP server connection is established, the ARP VPI/VCI is nonzero. The CALLING_SERVER flag indicates whether or not an attempt is currently being made to connect to the ARP server. The retry count indicates the number of attempts to connect to the ARP server which have been unsuccessful since the last successful attempt. The Inarp count indicates the number of InARP requests that have been sent on the ARP server connection without a response.

2.6 Configuring the Network

In an ATM network, before any connections can be made, the two parties must know each other's NSAP address in order to set up that connection.

To allow those connections to work, the ideal scenario is for all hosts and switches in the network to have support for both ILMI and for RFC-1577 (Classical IP over ATM). However, when using other vendors' equipment, this may not be the case. This section describes how to configure a network with the following scenarios:

- Configuring a third-party host that has no ILMI and no RFC-1577 support
- Configuring a third-party switch that has ILMI support, but no RFC-1577 support
- Configuring a third-party switch that has no ILMI support, but has RFC-1577 support

2.6.1 Third-Party Host with No ILMI and No RFC-1577 Support

To configure a network with a third-party vendor's host (or an edge device) that supports neither ILMI nor RFC-1577 (as shown in Figure 2.1), perform the following steps:

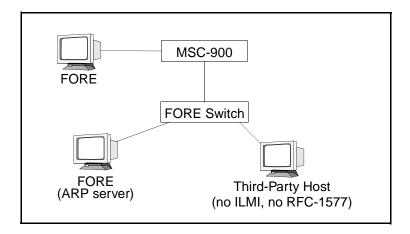


Figure 2.1 - Configuring a Third-Party Host with No ILMI and No RFC-1577 Support

- 1. Before beginning this process, be sure that FORE software is installed and running on the host that is the ARP server and on all of the other equipment.
- 2. Using the configuration software of the third-party host, assign that host an NSAP address that has the same prefix as the switch fabric to which it is connected.
- 3. Configure the switch so that it has a static route to the third-party host using the following AMI command:

conf nsap route new <NSAP> <mask> -port <port> -vpi <vpi>

Be sure to use a host mask value of 152.

4. Configure the ARP server (edit the firmware download script) with a static IP to NSAP mapping for the third-party host.

2.6.2 Third-Party Switch with ILMI and No RFC-1577 Support

To configure a network with a third-party vendor's switch that supports ILMI, but not RFC-1577, (as shown in Figure 2.2), perform the following steps:

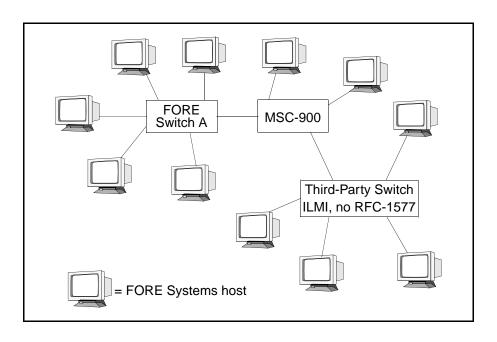


Figure 2.2 - Configuring a Third-Party Switch with ILMI Support and No RFC-1577

- Be sure that FORE software has been installed on all of the hosts and that ILMI was set in the process. Let ILMI dynamically perform address registration for all of the hosts.
- 2. Configure a static NSAP route to the third-party switch on the MSC-900 that is physically connected to the third-party switch using the following AMI command:

conf nsap route new <NSAP> <mask> -port <port> -vpi <vpi>

Be sure to use a network mask value of 104.

3. Configure two static NSAP routes on the third-party switch, one to each of the FORE devices to which the third-party switch is connected, using the third-party vendor's configuration software.

2.6.3 Third-Party Switch with RFC-1577 and No ILMI Support

To configure a network with a third-party vendor's switch that does not support ILMI, but does support RFC-1577 (as shown in Figure 2.3), perform the following steps:

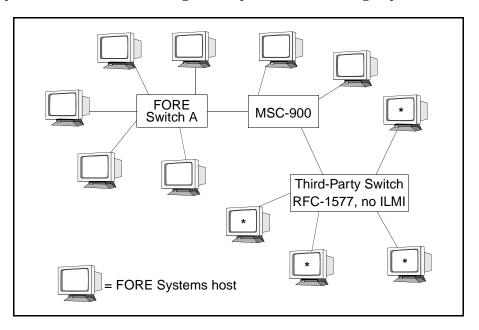


Figure 2.3 - Configuring a Third-Party Switch with RFC-1577 and No ILMI Support

- Be sure that FORE software has been installed on all of the FORE hosts and that ILMI was set in the process. Let ILMI dynamically perform address registration for all of the FORE hosts and FORE devices.
- 2. Statically configure the non-FORE (*) hosts with NSAP addresses (edit the firmware download script), using the same switch prefix for all of the hosts.
- Configure a static NSAP route to the third-party switch on the MSC-900 that is physically connected to the third-party switch using the following AMI command:

```
conf nsap route new <NSAP> <mask> -port <port> -vpi <vpi>
```

Be sure to use a network mask value of 104. Also, be sure to use the same prefix that was used to configure the hosts.

 Configure two static NSAP routes on the third-party switch, one to each of the FORE devices to which the third-party switch is connected, using the third-party vendor's configuration software.

CHAPTER 3

Configuring an Emulated LAN

3.1 Introduction

This chapter describes how to design, configure, and maintain an Emulated LAN (ELAN) over an ATM network. An ELAN provides communication of user data frames among all members of the ELAN, similar to a physical LAN. One or more ELANs may run simultaneously (and independently) on the same ATM network. Just as with physical LANs, communication between ELANs is possible only through routers or bridges.

3.1.1 Ethernet ELANs

The current software release supports emulation of Ethernet (IEEE 802.3) ELANs. Each ELAN is composed of a set of LAN Emulation Clients (LECs), a LAN Emulation Configuration Server (LECS), a LAN Emulation Server (LES), and a Broadcast and Unknown Server (BUS). Each LEC resides on an ATM host system (PC, Macintosh, UNIX workstation, switch, or bridge device). In the current software release, the LECS, LES, and BUS may reside either in an MSC-900 or in a UNIX workstation running SunOS version 4.1.x or Solaris 2.4 or 2.5. Additional software features include a colocated BUS (also referred to as an intelligent BUS or a LES/BUS pair) and a LEC "failover" mechanism that can provide redundancy to an ELAN.

3.1.2 Token Ring ELANs

The current software release supports emulation of Token Ring (IEEE 802.5) ELAN services. However, Token Ring LECs can not be created on the switch. In the current software release, the LECS, LES, and BUS may reside either in an MSC-900 or in a UNIX workstation running SunOS version 4.1.x or Solaris 2.4 or 2.5. Additional software features include a colocated BUS (also referred to as an intelligent BUS or a LES/BUS pair) and a "failover" mechanism that can provide redundancy to an ELAN.

3.2 Emulated LAN Components

The components of an emulated LAN include LECs, and LAN Emulation Services consisting of: a LECS, a LES, and a BUS. Each of these services may reside in the same physical system or in separate physical systems. For example, the LECS could reside in an MSC-900, while the LES and BUS reside in a workstation. In the current *ForeThought*-SP 1.0.x software, the LECS, LES, and BUS are supported only for MSC-900s and for systems running SunOS or Solaris. The functional interconnections of a simple ELAN consisting of two LECs, an LECS, a LES, and a BUS are shown in Figure 3.1.

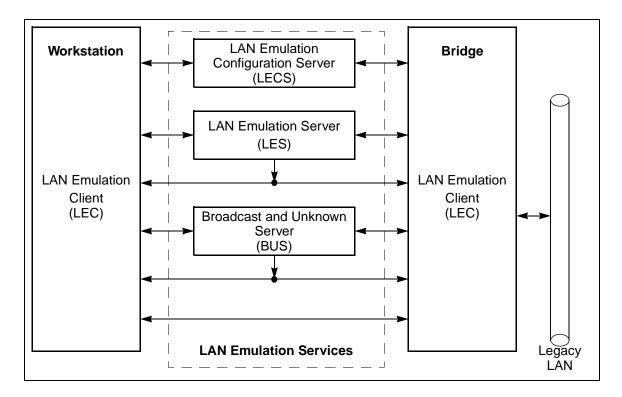


Figure 3.1 - Basic Emulated LAN Interconnections

Configuring an Emulated LAN

3.2.1 LAN Emulation Client (LEC)

The LEC is the component in an end system that performs data forwarding, address resolution, and other control functions when communicating with other components within the ELAN. It also provides a MAC level emulated Ethernet or Token Ring interface and appears to higher level software as though a physical interface is present. Each LEC must register with both the LES and BUS associated with the ELAN it wishes to join before it may participate in the ELAN. *ForeThought-SP* 1.0.x supports a maximum of 16 LECs per adapter card.

3.2.2 LAN Emulation Configuration Server (LECS)

The LECS is responsible for the initial configuration of LECs. It provides information about available ELANs that a LEC may join, together with the address of the LES associated with each ELAN. With *ForeThought-SP* 1.0.x, the user may also use the LECS to associate multiple LESs with a given ELAN. This feature allows LECs to "failover" to a hierarchy of redundant services if the primary LES for an ELAN goes down.

3.2.3 LAN Emulation Server (LES)

The LES implements the control coordination function for the ELAN. The LES provides the service of registering and resolving MAC addresses to ATM addresses. A LEC registers its own address with the LES. A LEC also queries the LES when the client wishes to resolve a MAC address to an ATM address. The LES either responds directly to the client or forwards the query to other clients so they may respond. There is only one instance of an active LES per ELAN.

3.2.4 Broadcast and Unknown Server (BUS)

Unlike traditional shared-media LAN architectures such as Ethernet or Token Ring, ATM is connection based. Therefore, it has no built-in mechanism for handling connectionless traffic such as broadcasts, multicasts, and unknown unicasts. In an ELAN, the BUS is responsible for servicing these traffic types by accepting broadcast, multicast, and unknown unicast packets from the LECs to the broadcast MAC address (FF-FF-FF-FF-FF) via dedicated point-to-point connections, and forwarding the packets to all of the members of the ELAN using a single point-to-multipoint connection. Each LEC is associated with only one active BUS. *Fore-Thought-SP* 1.0.x also supports the use of a colocated BUS (also referred to as an intelligent BUS or a LES/BUS pair) that allows the BUS to use the LES's registration table to direct unicast traffic.

3.3 Emulated LAN Operation

This section describes the operation of an ELAN and its components from the point of view of a LEC. The operation of an ELAN may be divided into three phases:

- 1. Initialization
- 2. Registration and Address Resolution
- 3. Data Transfer

ELAN components communicate with each other using ATM connections. LECs maintain separate connections for traffic control functions and data transfer. The following connection types are used by the LEC when operating in an ELAN:

- Configuration-Direct Connection: a bidirectional point-to-point VCC set up by the LEC to the LECS.
- *Control-Direct Connection:* a bidirectional point-to-point VCC set up by the LEC to the LES. This connection must be maintained for the duration of the LEC's participation in the ELAN.
- *Control-Distribute Connection*: a unidirectional point-to-multipoint VCC set up by the LES to the LEC. This connection must be maintained for the duration of the LEC's participation in the ELAN.
- *Multicast-Send Connection*: a bidirectional point-to-point VCC set up by the LEC to the BUS for sending multicast data to the BUS. The LEC must attempt to maintain this connection while participating in the ELAN.
- *Multicast-Forward Connection*: a unidirectional point-to-multipoint VCC set up from the BUS to LECs participating in the ELAN. This VCC must be established before a LEC participates in an ELAN. The LEC must attempt to maintain this connection while participating in the ELAN.
- *Data-Direct Connection*: a bidirectional point-to-point VCC set up between LECs that want to exchange unicast data traffic.

For the following discussion, please refer to Figure 3.2.

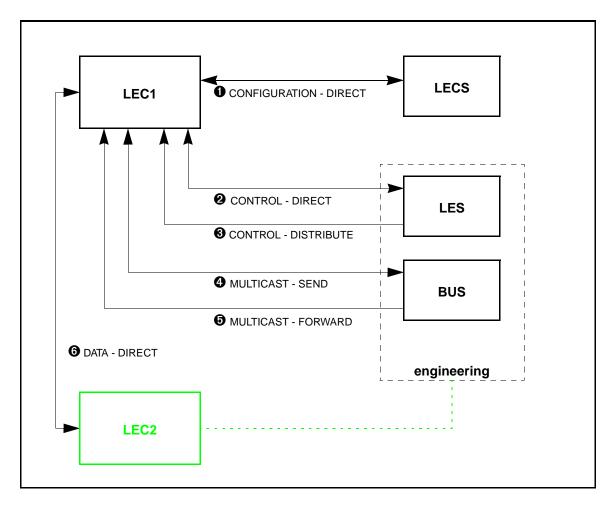


Figure 3.2 - ELAN Operation

3.3.1 Initialization

Once it knows the location of the LECS, LEC1 establishes a configuration-direct connection **1** to the LECS. When connected, the LECS provides LEC1 with the information necessary to connect to the ELAN it wishes to join. This information includes such parameters as: the ATM address of the ELAN's LES, the type of LAN being emulated, the maximum packet size, and the name of the ELAN (engineering, for example). This configuration information is contained in a configuration file that must be built and maintained by the network administrator.



Detailed information about the LECS configuration file may be found in Section 3.4.1.

3.3.2 Registration and Address Resolution

After obtaining the address of the LES, LEC1 establishes a control-direct connection **2** to the LES. The LES assigns LEC1 a unique identifier, and LEC1 registers its own MAC and ATM addresses with the LES. (The LES maintains a table containing the MAC addresses and corresponding ATM addresses of all members of the ELAN.) At this point, LEC1 has "joined" the ELAN.

The LES then establishes a control-distribute connection **3** back to LEC1. Connections **2** and **3** can now be used by LEC1 to send LAN Emulation ARP (LE_ARP) requests to the LES, and receive replies.

LEC1 now sends an LE_ARP request to the LES to get the ATM address of the BUS corresponding to the broadcast MAC address (FF-FF-FF-FF-FF). The LEC then establishes a multicast-send connection **4** to the BUS. The BUS responds by setting up a multicast-forward connection **5** to the LEC.

At this point, the LEC is ready to transfer data.

3.3.3 Data Transfer

When LEC1 receives a network-layer packet from a higher layer protocol to transmit to some destination MAC address (for example, LEC2), LEC1 initially does not know the corresponding ATM address of the destination. Consequently, LEC1 transmits an LE_ARP request to the LES.



The example shown in Figure 3.2 assumes that LEC2 has already registered with the LES, and that connections similar to those described for LEC1 already exist.

While waiting for the LES to respond, LEC1 forwards the packet to the BUS. The BUS broadcasts the packet to all LECs on the ELAN. This is done to avoid data loss, and to minimize connection set-up latency (due to the LE_ARP process) that may not be acceptable to some network protocols.

If the LE_ARP response is received, LEC1 establishes a data-direct connection **6** to the destination address of LEC2. This path will be used for subsequent data transfers. Before LEC1 begins to use this connection, it first sends a "flush" packet via the BUS to the destination, LEC2. When LEC2 acknowledges receipt of this packet, signifying that the BUS path is empty, only then does LEC1 begin to use the data-direct connection **6** for data transfer. This process ensures that the network protocol's frames arrive in the proper order.

If no response is received to the LE_ARP, LEC1 continues to send data via the BUS, while continuing to LE_ARP until a response is received and a data-direct connection to LEC2 is established.

If LEC1 already has a data-direct connection to a MAC address it wishes to reach, it need not go through the LE_ARP process again. Instead, it continues to use the current connection. This is possible because each LEC maintains a cache of MAC address to ATM address mappings that it receives in response to the LE_ARPs it has sent. Entries in this cache are "aged" out over a period of time. Data-direct connections are also cleared if they remain inactive for a period of time.

3.4 Configuring an ELAN

To configure an ELAN on an MSC-900, you must log into AMI on an MSC-900 running *Fore-Thought-SP* 1.0 or greater and use the commands found under configuration lane.



More information about each of these commands may be found in the *ForeThought-SP AMI Reference Manual*.

There are three major steps that the system administrator should follow in order to configure and maintain ELANs:

- 1. Configure an LECS configuration database file.
- 2. Start the LAN Emulation Services (LECS, LES, and BUS).
- 3. Start the LEC(s) and join an ELAN.



Steps 1 may be performed using a text editor on any system. However, for *ForeThought*-SP 1.0, the resulting file can be **used** only on systems running under SunOS Version 4.1.x, or Solaris 2.4 or 2.5, or on an MSC-900.

The remainder of this section gives a practical example of configuring and administering an ELAN using *ForeThought-SP* 1.0.x.

3.4.1 Configuring an LECS Configuration Database File

The LECS uses a text configuration file to contain the configuration information needed by LECs that wish to participate in an ELAN. The LECS configuration file may be built and edited using a text editor such as vi or emacs.



VLAN Manager, a product available separately from FORE Systems, greatly simplifies the administration of ELANs. Its graphical user interface controls the content of the LECS configuration file transparently to the user. For more information, please refer to the *ForeView VLAN Manager User's Manual*.

3.4.1.1 Before You Begin

Before building or modifying the LECS configuration file, you should first determine the topology of the ELAN or ELANs that you want to administer. You must supply the following information when building or editing the LECS configuration file:

- Provide the name of each ELAN (engineering, marketing, etc.).
- Provide the LAN type (Ethernet or Token Ring) for each ELAN.
- Provide the MTU size for each ELAN.
- Provide the ATM address of the LES for each ELAN.
- If LEC failover is to be made available to clients on a given ELAN, the ATM addresses for all LESs in that ELAN's failover sequence must be included.
- Provide the address of each LEC that may participate in each ELAN.
- If you wish LECs to use a default ELAN, the default LES information must also be included.
- Provide various other configurable parameters.

CAUTION



Do not attempt to edit an existing functional LECS configuration file without first making a backup copy of the file. Incorrect modification of the configuration file could result in loss of communication between one or more members of a defined ELAN.



You may make changes to the LECS configuration file while the LECS process is running. The configuration file is reread periodically by the LECS process (the default period is ten minutes). Consequently, any changes that you make to the configuration file are not recognized until the file is reread.

3.4.1.2 LECS Configuration File Syntax

Each line that you enter in the configuration file takes the general form:

[[group].]key : value

The group field may represent:

- ELANs (by name) ELAN names are case-sensitive, and may not exceed 32 characters in length
- clients ATM or MAC addresses
- miscellaneous LECS control information specified by using a group name of LECS

The key field is used to denote an individual parameter within a group.

The value field contains the value assigned to the key.

Omitting the group implies that the key and value apply to all groups in the configuration file. Leading and trailing spaces, as well as spaces on either side of the ":", are ignored.

For example, to specify a maximum frame size of 1516 bytes for the ELAN named engineering, enter the following:

engineering.Maximum Frame Size: 1516

Similarly, to specify a default maximum frame size of 1516 bytes for all ELANs defined in a given configuration file, enter the following:

.Maximum_Frame_Size : 1516

Table 3.1 defines the various key parameters that may be entered in the configuration file. The default value for each parameter is also given.

Table 3.1 - LECS Configuration File Parameters

Parameter	Definition
.LAN_Type: Ethernet/IEEE 802.3	Identifies the type of ELAN, either Ethernet/IEEE 802.3 or IEEE 802.5. The default is Ethernet/IEEE 802.3.
.Maximum_Frame_Size: 1516	Specifies the length (in number of bytes) of the largest frame field. Selections are: 1516, 4544, 9234, and 18190. The default is 1516 for Ethernet and 4544 for Token Ring.
.Control_TimeOut: 120	Specifies the timing out of request/response control frame interactions, in seconds. The default is 120 seconds.
.Maximum_Unknown_Frame_Count: 1	Limits the number of unicast frames sent to the BUS. The default is 1 frame.
.Maximum_Unknown_Frame_Time: 1	Limits the number of unicast frames sent to the BUS in the specified number of seconds. The default is 1 second.
.VCC_TimeOut_Period: 1200	Specifies the length of time that an idle data connection remains open before being closed. The default value is 1200 seconds.
.Maximum_Retry_Count: 1	Limits the number of LE_ARP requests. The default is 1.
.Aging_Time: 300	Specifies the period that LE_ARP cache table entries remain valid, in seconds. The default value is 300 seconds.
.Forward_Delay_Time: 15	Specifies the timing out of non-local ARP cache entries in seconds. The default value is 15 seconds.
.Expected_LE_ARP_Response_Time: 1	Specifies the maximum time a LEC expects an LE_ARP request/response will take, in seconds. The default value is 1 second.
.Flush_TimeOut: 4	Specifies the maximum time a LEC expects an LE_FLUSH request/response will take, in seconds. The default value is 4 seconds.
.Path_Switching_Delay: 6	Specifies the minimum time between switching BUS and data paths, in seconds. The default value is 6 seconds.

Parameter	Definition
.Multicast_Send_VCC: Best Effort	Specifies the multicast send mode, either Best Effort, Variable, or Constant. The default is Best Effort.
.Connection_Complete_Timer: 4	Specifies the time period in which data or READY_IND is expected, in seconds. The default is 4 seconds.

Table 3.1 - LECS Configuration File Parameters (Continued)

Lines beginning with # may be inserted if you wish to include comments or to improve the clarity of the presentation when the file is viewed or printed. These lines are ignored when the file is read. Lines may be continued by escaping the end-of-line with a backslash "\" (do not enter the quote marks).

3.4.1.3 Defining an ELAN

Each ELAN is defined by an address statement whose value denotes the ATM address of the ELAN's LES. For example:

```
engineering.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.11
```

If you wish to define multiple instances of an ELAN for use by the LEC failover mechanism, the address statements would appear as follows:

```
marketing | 0.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.21 marketing | 1.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.23 marketing | 2.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.25
```

In this example, redundant LES addresses for the marketing ELAN have been defined. Note that the group name, marketing, ends with the meta-character |, followed by a number indicating the position of the LES in the failover sequence. If the marketing | 0 LES were to fail, LECs could "failover" to the marketing | 1 LES. The LEC then periodically tries to reconnect with marketing | 0. If marketing | 1 also failed, the LEC would then failover to marketing | 2, (and periodically attempt to reconnect to marketing | 0 and marketing | 1).

In addition, you may instruct a given ELAN to override any of the default values. For example, the engineering ELAN could override a Maximum_Frame_Size setting of 1516; thus:

engineering.Maximum Frame Size: 4544

If you want to control which clients may or may not join a given ELAN, two additional keys, Accept and Reject, whose values are comma-separated lists of matching elements, may be used. These values may be:

a MAC address,

engineering.Accept: 0020480605b2 , 002048080011 , 0020481020ef an ATM address and equal-length bit mask,

engineering.Accept: 47.0005.80.FFE100.0000.0000.0000.002048000000.00\ FF.FFF.FF.FFFFF.0000.0000.FFFFFF000000.00

or an ATM address containing "don't-care" semi-octets denoted by an "X".

marketing.Accept: 47.0005.80.FFE100.XXXX.XXXX.XXXX.002048XXXXXX.XX

The last two forms of ATM-address matching elements are functionally the same. The latter is shorter but only allows for masks whose semi-octets are all ones or all zeros, while the former allows for arbitrary masks. A prospective-client address is "captured" by an ELAN name if the client's address matches one of the Accept elements but not one of the Reject elements (if present). Finally, an ELAN may be configured to accept any client that wishes to join by including the following statement:

The order in which to apply the Accept and Reject rules is given by a Match.Ordering group.key statement, whose value is a comma-separated list of ELAN names. For example:

Match.Ordering: default, engineering, marketing 0, marketing 1, marketing 2

The names of all ELANs that have Accept keys must be included in Match. Ordering. Names of all of the members of a LEC failover sequence must be included so that LECs may use them.

3.4.1.4 Defining a Client

Clients need not be defined in the LECS configuration file. Typically, you would define a client for the purpose of overriding one or more of the default configuration parameters for that particular client.

A client is defined by using its ATM address, MAC address, or route descriptor in the group field, and perhaps giving the name of its ELAN as the value of the LAN_Name key. For example:

47.0005.80.FFE100.0000.0000.0000.002048222222.22.LAN Name:engineering

002048ABCDEF.LAN Name: marketing

012F.LAN Name: publications

Configuration parameter overrides can also be given on a per-client basis. For example, the following statements override the default VCC_TimeOut_Period and Aging_Time configuration parameters for a client whose MAC address is 002048080011 on the engineering ELAN:

002048080011.LAN Name:engineering

002048080011.VCC TimeOut Period:1200

002048080011.Aging_Time: 30

3.4.1.5 LECS Control Parameters

Specifying values for keys in the LECS group provides control over the operation of the LECS process.



If you change the values of the LECS control parameters while the LECS process is running, the new values do not take effect until the LECS process is stopped, and then restarted.

When a client contacts the LECS, the connections established are known as Configuration Direct VCCs. To override the default value of the VCC_TimeOut_Period key (the number of seconds before an idle Configuration Direct VCC is automatically closed by the LECS), enter a statement similar to the following:

LECS.VCC TimeOut Period: 1200

The LECS periodically checks whether its configuration file has been modified, and, if it has, the file is reread. The length of this period, in seconds, is given by the Reload_Period key:

LECS.Reload Period: 600

The Permanent_Circuits key holds a comma-separated list of VPI.VCI pairs denoting the local ends of 0.17 PVCs on which the LECS should listen. For example:

LECS.Permanent Circuits: 0.42, 0.112

The LECS can provide the client with a fourteen-bit pattern to permute the MAC-address generation algorithm. This bit pattern is specified with the MAC_Address_Base key.

LECS.MAC Address Base: 38fe

3.4.2 Sample LECS Configuration File

CAUTION



Do not attempt to edit an existing functional LECS configuration file without first making a backup copy of the file. Incorrect modification of the configuration file could result in loss of communication between one or more members of a defined ELAN.



For a detailed discussion of how to configure an LECS configuration file similar to the one given in this section, please refer to Section 3.4.1.

The sample LECS configuration file shown at the end of this section in Figure 3.3 and Figure 3.4 defines three ELANs:

- default
- engineering
- marketing

The Match.Ordering statement specifies the ELAN names in the order that prospective clients will attempt to join. Note that redundant marketing ELANs have been included. This allows clients attempting to join the marketing ELAN to "failover" to one of the redundant LES instances if the primary LES fails.

The default configuration parameters are shown with their default values. These values apply to all ELANs in this configuration file, unless overridden for a particular ELAN or client.

ELAN default is configured to accept any client that wishes to join. The administrator must substitute the ATM address of the default LES in place of the "4"s in the default.Address statement.

ELAN engineering has overridden the default Maximum_Frame_Size with a new size of 4544 bytes. Consequently, this frame size applies only to traffic on the engineering ELAN. The default and marketing ELANs continue to use the default frame size of 1516 bytes.

ELAN marketing has three instances, |0,|1, and |2. These three instances provide LECs that join the ELAN marketing the information necessary to implement the LEC "failover" mechanism. marketing |0| defines the primary services for the ELAN. If marketing |0| fails, LECs connected to it can automatically switch to marketing |1| and continue operating (while periodically attempting to rejoin marketing |0|). If marketing |1| were also to fail, the LECs could then switch to services on marketing |2|.

Two LECs, whose MAC addresses are 002048080011 and 0020481020ef, are identified as acceptable clients for the engineering and marketing ELANs.

```
# The search ordering of elan names
Match.Ordering: default, engineering, marketing 0, marketing 1, marketing 2
# the default configuration parameters
  .Control_TimeOut: 120
 .Maximum_Unknown_Frame_Count: 1
 .Maximum_Unknown_Frame_Time: 1
 .VCC_TimeOut_Period: 1200
 .Maximum_Retry_Count: 1
 .Aging_Time: 300
 .Forward_Delay_Time: 15
  .Expected_LE_ARP_Response_Time: 1
 .Flush TimeOut: 4
  .Path_Switching_Delay: 6
 .Multicast_Send_VCC_Type: Best Effort
  .Connection_Complete_Timer: 4
 .LAN_Type: Ethernet/IEEE 802.3
  .Maximum_Frame_Size: 1516
 # Parameters for the active default elan
```

Figure 3.3 - Sample LECS Configuration File (Part One of Two)

```
# Parameters for elan: engineering
# engineering.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.11
engineering.Accept: 002048080011 , 0020481020ef
engineering.Maximum_Frame_Size: 4544
#
# Parameters for elan: marketing|0
#
marketing|0.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.21
marketing|0.Accept: 002048080011 , 0020481020ef
#
# Parameters for elan: marketing|1
#
marketing|1.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.23
marketing|1.Accept: 002048080011 , 0020481020ef
#
# Parameters for elan: marketing|2
#
marketing|2.Address: 47.0005.80.ffe100.0000.f21a.01b9.0020480605b2.25
marketing|2.Accept: 002048080011 , 0020481020ef
```

Figure 3.4 - Sample LECS Configuration File (Part Two of Two)

3.4.3 The Default LECS Configuration File

A default LECS configuration file is installed with *ForeThought-SP* 1.0.x. It is reproduced below for reference.

Figure 3.5 - Default LECS Configuration File



Make a backup copy of the supplied LECS configuration file and work with the copy.

As supplied, this file allows you to set up a default ELAN that accepts any client that wants to join. Before you can use the supplied file, you must modify it as follows:

- Obtain the ATM address of the machine where the LES for the default ELAN will
 be started. You must then modify the machine's ATM address by changing the
 selector byte. This modified address is the same address that is used when starting
 the default LES.
- 2. Substitute the address determined in step 1 for the string of "4"s in the default.Address statement of the LECS configuration file. The default.Accept string of "X"s should not be changed as this parameter allows any client who wishes to join the default ELAN.
- 3. Save the modified file.



For a detailed discussion of how to configure an LECS configuration file other than the supplied default file, please refer to Section 3.4.1.

3.4.4 Starting the LAN Emulation Services

LAN Emulation services include the LECS, LES, and BUS. Once the LECS configuration database file has been configured, each of these services must be started so that they are available for LECs to attempt to use. These services may, but need not, run in the same machine.

3.4.4.1 Starting the LECS

Once an LECS configuration file has been configured, you need to retrieve the LECS configuration database file that you built by using the following AMI command:

```
configuration lane lecs get <host>:<remotefile> [<localfile>]
```

For example, you would enter something similar to the following:

configuration lane lecs get 198.29.22.46:lecs.cfg



The default local file for an MSC-900 is lecs.cfg.

On an MSC-900, this file is retrieved via tftp.

After you have retrieved the LECS configuration database file, use the following AMI command to start the LECS service on the MSC-900:

For example, to start the LECS service using the -db option and using a well-known address other than the ATM Forum's well-known address, you would enter something similar to the following:

```
configuration lane lecs new 0x0c -db lecs.cfg 0x47.0005.80.ffe100.0000.f215.11f1.0020481511f8.99
```

Use the following AMI command to verify that the LECS has been started and is running.

configuration lane lecs show

Index	AdminStatus	OperStatus	Selector	WKA	Database
1	up	up	0x0c	other	lecs.cfg
	wka: 0x47.00	05.80.ffe100.	0000.f215.	11f1.0020481	511f8.99

The OperStatus field displays up, meaning that the LECS is enabled. Because an ATM address has been entered at the end of the command, the WKA field shows other, meaning that an address other than the ATM Forum's well-known address is being used. This address is displayed below the entry. Because the <code>-default</code> option has not been specified, the default LES field is not displayed. If you use the <code>-default</code> option, then that LES address is displayed below this entry in a field titled <code>Default LES</code>.

3.4.4.2 Creating a LES and a BUS

Next, the LES and BUS services must be started for the ELAN. The preferred method is to use the colocated_bus option when creating the LES, which means that the LES and BUS services for a particular ELAN will be started together and will be running on the same MSC-900. This configuration may provide better BUS performance. To do this, use the following AMI command:

For example, you would enter something similar to the following:

conf lane les new a b default colocated_bus



By using the colocated_bus option, you are creating a LES and BUS using a single AMI command. There is no need to create a BUS separately.



If no type is entered, the MSC-900 assumes Ethernet and uses 1516 as the MTU size. If Token Ring is used as the type, but no MTU size is entered, 4544 is used as the size.

Use the following AMI command to verify that the LES and the BUS have been started and are running. The <code>OperStatus</code> field in this listing should display up, meaning that the LES and the BUS are enabled.

configuration lane les show

3.4.5 Starting the LEC(s) and Joining an ELAN

Now that the ELAN services have been started, you can have LECs join the ELAN that you have created.



The MSC-900 software only allows you to create an instance of a LEC on an MSC-900. To create an instance of a LEC on a host, you must use the VLAN Manager or use a FORE host adapter. Please refer to the respective User's Manual for instructions.



A LEC created on the MSC-900 cannot join a Token Ring ELAN. It can only join an Ethernet ELAN.

3.4.5.1 Creating a LEC

To start a LEC that will attempt to join the ELAN, use the following AMI command:



The recommended method for starting a LEC is to use the automatic mode, meaning that the LEC will attempt to contact the LECS on the "well-known" address as defined by the ATM Forum's LAN Emulation standards (47.0079.00.0000000.0000.0000.0000.000 A03E000001.00).

For example, to start a LEC that attempts to join the default ELAN called default (assuming that a default ELAN has been defined in the LECS configuration file), enter the following:

configuration lane lec new 0 default

If you decide to use the manual mode, you must enter either a LECS address other than the well-known address or you must enter a LES address. If you enter a LES address, this means that the LEC bypasses the LECS and directly contacts the specified LES.

configuration lane lec new aa default manual -les 47000580ffe1000000f21a00d50020481a00d50c

Regardless of which way you start the LEC, you can verify that the LEC has joined the ELAN by using the following AMI command:

configuration lane lec show

	Admin	Oper					
Index	Status	Status	Sel	Mode	MACaddress	IfName	ELAN
1	up	up	0x00	automatic	000000000000	el0	default
	LECS: 0x4	47.0079.0	0.000	000.0000.000	0.0000.00a03e00	00001.00	
	LES : 0x4	47.0005.8	30.ffe1	100.0000.f21a	a.00d0.0020481a	06.0b00e	

Because the OperStatus field is displaying joining, this means that the LEC is still registering with the ELAN. When it has finished, the OperStatus field displays up.

After the first LEC has joined the ELAN, you can use the VLAN Manager or the host software to add more LECs to this ELAN. Once all the LECs have joined, the ELAN is complete.

3.4.5.2 Configuring the LEC Failover Mechanism

An added feature of *ForeThought-SP* 1.0.x software is the LEC failover mechanism. This feature allows multiple redundant services to be configured for an ELAN as shown in the sample LECS configuration file in Section 3.4.2. Consequently, to start LECs that join ELAN marketing and make use of the failover mechanism set up in the sample LECS configuration file, enter the following AMI commands:

configuration lane lec new aa marketing | 0

followed by,

configuration lane lec new ab marketing 1

and then,

configuration lane lec new ac marketing 2

The above commands start three LECs. Initially, the LEC that connects to ELAN marketing | 0 services is the active LEC. If marketing | 0 services fail, the LEC associated with ELAN marketing | 1 services becomes active. Periodically, the LEC associated with ELAN marketing | 0 attempts to re-establish service. If both marketing | 0 and marketing | 1 services fail, the LEC associated with ELAN marketing | 2 then becomes active. You may define as much redundancy into your system as resources allow.

CAUTION



Deleting (stopping) the currently active failover LEC removes the network interface (e.g., el0) associated with **all** of the failover LECs, causing the other LECs in the failover mechanism to lose IP connectivity.

CHAPTER 4

Configuring Circuit Emulation Services

FORE Systems' Circuit Emulation Services (CES) Network Modules (NMCE-6/DS1A and NMCE-6/E1A) provide adaptation from time-division multiplexed (TDM) equipment (i.e., PBXs, WAN multiplexers, channel banks, video codecs, etc.) and traffic to ATM. Both modules provide structured and unstructured services, with a maximum of 127 connections supported on each module.

- All six ports of the DS1 CES Network Module may support fractional DS1 services (n x 56 Kbps/n x 64 Kbps) where 1 to 24 contiguous or non-contiguous DS0 channels are mapped to a single ATM VCC not to exceed 127 total connections.
- All six ports of the E1 CES Network Module may support fractional E1 services (n x 56 Kbps/n x 64 Kbps) where 1 to 31 contiguous or non-contiguous DS0 channels are mapped to a single ATM VCC not to exceed 127 total connections.

Structured services provide digital access and cross-connect system (DACS) connectivity where n x 64 Kbps and n x 56 Kbps digital signal level zero (DS0) channels are adapted to ATM cells and mapped to unique ATM virtual connections (VCCs).

Unstructured services provide support and maintenance of a single full bandwidth 1.544 Mbps (DS1) or 2.048 Mbps (E1) clear channel across a single ATM virtual connection.

Configurations of both the DS1 and the E1 version of the CES network module are detailed in this chapter. This chapter assumes the proper configuration of general MSC-900 parameters and of CES specific parameters, such as timing distribution/recovery, additional CES alarms, CES ports, etc.

4.1 Configuring CES Connections

The ces commands let the user create and delete CES connections, as well as display the status of existing connections. You can display the list of available subcommands by typing ces ? at the configuration level.

4.1.1 Creating a New CES Connection

To create a new CES connection, the user must set several parameters. Enter the following to create a new CES connection:

```
localhost::configuration ces> new <port> <timeslots>
```

The CES new command can also be used as shown below. When the following parameters are used, by default, an appropriate entry is made in the UPC table and a bidirectional PVC is created with the proper UPC index.

```
new <port> <timeslots> -oport <oport> -ots <ots>
new <port> <timeslots> -oport <oport> -ovpi <ovpi> -ovci <ovci>
```

The following advanced options can be used when creating CES connections:

```
[-srts (on|off)] [-fupc <index>] [-bupc <index>]
[-cas (basic|cas)] [-partialfill <partialfill>] [-reasscdvt <cdvt>]
[-bufSize <bufSize>] [-integ <integ>]
```



SRTS is only available on unstructured connections, which are created by specifying all for the <timeslots> parameter.

The -cas and -partialfill options are not applicable to unstructured mode.

Structured mode is selected by indicating the exact timeslots to be used. For example, timeslots 1, 2, and 3 would be entered as 1-3, timeslots 2, 4, and 6 would be entered as 2:4:6, and combinations such as 1-4:9-11:12 are allowed.

The parameters in the **new** command have the following meanings:

port Indicates the network module port on which the CES connection is to be created.

timeslots Indicates which timeslots (1-24 for DS1, 1-31 for E1) are being configured for a particular PVC. all indicates unstructured service. The time slot assignments may be either contiguous or noncontiguous DS0s.

oport Indicates the output port of the CES connection, which can be a CES port or an ATM port.

ovpi Indicates the output Virtual Path Identifier (VPI) of the CES connection when the output port is not a CES port.

ovci Indicates the output Virtual Channel Identifier (VCI) of the CES connection when the output port is not a CES port.

ots Indicates the output timeslots of the CES connection when the output port is a CES port.

srts Indicates whether Synchronous Residual Time Stamp (SRTS) clock recovery is to be enabled on this connection. on indicates that SRTS is enabled, off indicates that SRTS is disabled. By default, SRTS is disabled.

fupc Indicates the UPC contract type to be used in the ingress direction of the connection. (See Chapter 2 in the *ForeThought-SP AMI Reference Manual* for more information about UPC contracts.)

bupc Indicates the UPC contract type to be used in the egress direction of the connection. (See Chapter 2 in the *ForeThought-SP AMI Reference Manual* for more information about UPC contracts.)

cas Indicates whether Channel Associated Signalling (CAS) is to be used on the connection. basic, which applies to structured connections only, indicates that CAS will not be used, cas indicates that CAS will be used. The default is basic.

pa	rtı	2	ltı	ш

Indicates how many of the available 47 payload bytes in each cell are used before they are deemed "full" and ready for transmission across the ATM network (i.e. how much of the ATM cell contains data and how much is filler). The range for this parameter is 5 to 47. The default value is 47, for 47 bytes of data.

partialfill is used to minimize network transmission latency and is useful especially with time-sensitive, robbed-bit signalling sources.

reasscdvt <cdvt>

Indicates the Cell Delay Variation Tolerance for cells being received by the segmentation and reassembly (SAR) engine. The range for this parameter is 100 to 24000 (in μ s), and the default is 2000 (i.e., 2 ms).

bufSize

Indicates the amount of reassembly buffer space allocated for the connection. The default is 512 bytes per timeslot.

integ

Indicates the amount of time allocated to re-establish the connection before, while, or after the call is established, or in the case of interruption. The default value is $2500\mu s$.

4.1.2 Displaying CES Connections

To display the current CES connections, enter the following:

localhost::configuration ces> show

CES		Input					Outr	put		
Service	State	Port	Timeslots	VPI	VCI	Type	Port	TimeSlots	VPI	VCI
24	down	1A1	1	0	129	-	-	_	-	-
31	down	1A1	2-3	0	130	spvc	1D3	_	0	32
32	down	1A1	4-5	0	131	pvc	1D4	_	0	150
33	down	1A1	6-7	0	132	spvc	1D3	=	0	35

The fields in this display are explained in Section 4.1.1, except for the following:

CES Service Indicates the identification number (assigned by the

MSC-900) of the CES connection.

State Indicates whether the CES connection is enabled (up)

or disabled (down).

Input VPI Indicates the incoming VPI value of the connection.

Input VCI Indicates the incoming VCI value of the connection.

Type Indicates the type of ATM connection (i.e., PVC or

SPVC) that is associated with the CES connection.

To display the advanced settings of the current CES connections, enter the following:

localhost::configuration ces> show advanced
Searching for CEM Ports...

			Service	Clock		Partial	Max		Integ.
Service	MapVPI	MapVCI	Type	Mode	Cas	Fill	BufSize	CDVT	Period
31	0	129	structured	synch	basic	47	256	2000	2500

The fields in this display are explained in Section 4.1.1, except for the following:

Service Indicates the identification number (assigned by the

MSC-900) of the CES connection.

MapVPI Indicates the incoming VPI value of the connection.

MapVCI Indicates the incoming VCI value of the connection.

CDVT See reasscdvt <cdvt> in Section 4.1.1.

Integ. Period See integ in Section 4.1.1.

Configuring Circuit Emulation Services

APPENDIX A Configuring SNMP

The software for the MSC-900 includes an SNMP agent. The SNMP agent enables the remote monitoring and configuration of these devices.

A.1 SNMP Indexing

There are two main SNMP indexing schemes used: software port indices and hardware port indices. Software port indices are single numbers starting at 0 for the first port, incrementing 8 ports per module on an MSC-900. For example, port A1 on an MSC-900 has a software port index of 0. Port C3 on an MSC-900 has a software port index of 18, or 8 * 2 + 2.

Hardware port indices are of the form {board}.{network module}.{port} or bnp notation. They start at 0.0.0 for the first port and increment across boards, network modules, and ports. For example, port C3 on an MSC-900 is hardware port 0.2.2.

Please refer to Table A.1 for a summary of the port numbering conventions used in the MSC-900 and related SNMP indexing format.

Table A.1 - MSC-900 Port Numbering

Port Name	Software Port Number	Board-Netmod- Port Index	Port Name	Software Port Number	Board-Netmod- Port Index
A1	0	0.0.0	C1	16	0.2.0
A2	1	0.0.1	C2	17	0.2.1
A3	2	0.0.2	C3	18	0.2.2
A4	3	0.0.3	C4	19	0.2.3
A5	4	0.0.4	C5	20	0.2.4
A6	5	0.0.5	C6	21	0.2.5
B1	8	0.1.0	D1	24	0.3.0
B2	9	0.1.1	D2	25	0.3.1
В3	10	0.1.2	D3	26	0.3.2
B4	11	0.1.3	D4	27	0.3.3
B5	12	0.1.4	D5	28	0.3.4
В6	13	0.1.5	D6	29	0.3.5
			CTL	56	0.7.0

A.2 SNMP Traps

SNMP traps are used to update the state of the network automatically to remote network management hosts. The SNMP agent on the MSC-900 supports several SNMP traps.

The traps generated by the MSC-900's SNMP agent can be sent to as many destinations as needed. These destinations are configurable via the ATM Management Interface (AMI). Each destination must be an IP address of a network management host. The network management host specified for a trap destination can be any host with which the MSC-900 has connectivity. This means that the host does not have to be a directly connected ATM host. It can be on any attached network. The following table describes the supported traps.

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
0	asxSwLinkDown	An asxSwLinkDown trap signifies that the sending protocol entity recognizes a failure in one of the ATM Switch links that is connected to another switch.
1	asxSwLinkUp	An asxSwLinkUp trap signifies that the sending protocol entity recognizes that one of the ATM Switch links that is connected to another switch has come up.
2	asxHostLinkDown	An asxHostLinkDown trap signifies that the sending protocol entity recognizes a failure in one of ATM Switch links that is connected to a host.
3	asxHostLinkUp	An asxHostLinkUp trap signifies that the sending protocol entity recognizes that one of the ATM Switch links that is connected to a host has come up.
4	asxNetModuleDown	An asxNetModuleDown trap signifies that the sending protocol entity recognizes a failure in one of the ATM Switch network modules, that is identified by the board and the module numbers. This is probably caused by a hot-swap of a network module.
5	asxNetModuleUp	An asxNetModuleUp trap signifies that the sending protocol entity recognizes a new operational ATM Switch network module, that is identified by the board and the module numbers. This is probably caused by a hot-swap of a network module.

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
6	asxPsInputDown	This trap alerts that one ATM switch power supply failed due to failure in the input voltage. The power supply that failed is identified by the power supply index. Note that an input voltage may be out of specification and may not cause a power supply failure if high loads are not applied.
7	asxPsInputUp	This trap alerts that one ATM switch power supply that had an AC input failure is up. The power supply that is back up is identified by the power supply index.
9	asxPsOutputDown	This trap alerts that one ATM switch power supply output or the power supply was physically removed. The power supply that failed is identified by the power supply index.
10	asxPsOutputUp	This trap alerts that one ATM switch power supply that had an output failure or was removed is now up. The power supply that is back up is identified by the power supply index.
22	asxFanBankDown	This trap alerts that one ATM switch fan bank failed. The fan bank that failed is identified by the fan bank index.
23	asxFanBankUp	This trap alerts that one ATM switch fan bank is up. The fan bank that is back up is identified by the fan bank index.
28	asxLinkDown	This trap alerts that the link that is identified by {hwPortBoard, hwPortModule, hwPortNumber} was configured up but lost its carrier (or the framing bit) and is currently down.
29	asxLinkUp	This trap alerts that the link that is identified by {hwPortBoard, hwPortModule, hwPortNumber} is back up.
30	asxSpansDown	This trap alerts that the SPANS signalling on the link that is identified by the sigPathPort and sigPathVPI failed.

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
31	asxSpansUp	This trap alerts that the SPANS signalling on the link that is identified by the sigPathPort and sigPathVPI is up.
32	asxTempSensorOverTemp	This trap alerts that one of the temperature sensors indicates over temperature. The temperature sensor is identified by the temperature sensor index.
33	asxTempSensorRegularTemp	This trap alerts that one of the temperature sensors indicates regular temperature. The temperature sensor is identified by the temperature sensor index.
34	asxFabricTemperature OverTemp	This trap alerts that one of the temperature sensors indicates over temperature. The temperature sensor is identified by the temperature sensor index.
35	asxFabricTemperature RegularTemp	This trap alerts that one of the temperature sensors indicates regular temperature. The temperature sensor is identified by the temperature sensor index.
36	asxSonetLOSon	This trap indicates that the specified SONET port is experiencing Loss Of Signal. Bellcore Document TA-NWT-000253 Section 6.3.1.1.1 states that A SONET NE shall declare a LOS failure when the LOS defect persists for 2.5 (+5) seconds, or when a LOS defect is present and the criteria for LOF failure declaration have been met.
37	asxSonetLOSoff	This trap indicates that the LOS condition identified by trap asxSonetLOSon has been cleared.
38	asxSonetPathLabelOn	This trap indicates that the specified SONET port is receiving and errored C2 Path Label byte. Reference Bellcore Document TA-NWT-000253 Section 3.3.2.3 and 6.3.1.1.8 the Path Label (C2) byte should have the value 0x13.
39	asxSonetPathLabelOff	This trap indicates that the Errored Path Label (C2) byte error condition signalled by the asxSonetPath-LabelOn trap has been cleared.

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
40	asxSonetLineAISon	This trap indicates that the specified SONET port is receiving a Line level Alarm Indication Signal from the far-end equipment.
41	asxSonetLineAISoff	This trap indicates that the Line AIS error condition signalled by the asxSonetLineAISon trap has been cleared.
42	asxDS3FERFOn	This trap indicates that the specified DS3 port is in the DS3 Yellow Alarm or FERF state. The FERF or DS3 Yellow alarm is declared if either OOF(LOF), LOS or AIS is detected and persists for 2.5+5 seconds.
43	asxDS3FERFOff	This trap indicates that the specified DS3 port is no longer in the FERF or DS3 Yellow Alarm state.
44	asxDS3PLCPYellowOn	This trap indicates that the specified DS3 port is in the PLCP Yellow Alarm state. The Yellow alarm is declared if PLCP LOF is detected and persists for 2.5+5 seconds.
45	asxDS3PLCPYellowOff	This trap indicates that the specified DS3 port is no longer in the PLCP Yellow Alarm state.
46	asxDS3PLCPYellowDetected	This trap indicates that the specified DS3 port has detected incoming Yellow Alarm.
47	asxDS3PLCPYellowCleared	This trap indicates that the specified DS3 port has detected clearance of incoming Yellow Alarm.
48	asxDS3PLCPLOFDetected	This trap indicates that the specified DS3 port has detected incoming LOF Alarm.
49	asxDS3PLCPLOFCleared	This trap indicates that the specified DS3 port has detected clearance of incoming LOF Alarm.
50	asxDS3LOFDetected	This trap indicates that Loss Of Frame(LOF) is detected on the incoming signal.
51	asxDS3LOFCleared	This trap indicates that Loss Of Frame is cleared on the incoming signal.
52	asxDS3AISDetected	This trap indicates that AIS Alarm is detected on the incoming signal.

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name Description					
53	asxDS3AISCleared	This trap indicates that AIS Alarm is cleared on the incoming signal.				
54	asxDS3LOSDetected	This trap indicates that LOS Alarm is detected on the incoming signal.				
55	asxDS3LOSCleared	This trap indicates that LOS Alarm is cleared on the incoming signal.				
56	asxDS1YellowOn	This trap indicates that the specified DS1 port is in the Yellow Alarm state. The Yellow alarm is declared if either OOF or AIS is detected and persists for 2.5+5 seconds.				
57	asxDS1YellowOff	This trap indicates that the specified DS1 port is no longer in the Yellow Alarm state.				
58	asxDS1PLCPYellowOn	This trap indicates that the specified DS1 port is in the PLCP Yellow Alarm state. The Yellow alarm is declared if PLCP LOF is detected and persists for 2.5+5 seconds.				
59	asxDS1PLCPYellowOff	This trap indicates that the specified DS1 port is no longer in the PLCP Yellow Alarm state.				
60	asxDS1PLCPYellowDetected	This trap indicates that the specified DS1 port has detected an incoming Yellow Alarm.				
61	asxDS1PLCPYellowCleared	This trap indicates that the specified DS1 port has detected clearance of an incoming Yellow Alarm.				
62	asxDS1PLCPLOFDetected	This trap indicates that the specified DS1 port has detected an incoming LOF Alarm.				
63	asxDS1PLCPLOFCleared	This trap indicates that the specified DS1 port has detected clearance of an incoming LOF Alarm.				
64	asxDS1YellowDetected	This trap indicates that Yellow Alarm is detected on the incoming signal.				
65	asxDS1YellowCleared	This trap indicates that Yellow Alarm is cleared on the incoming signal.				
66	asxDS1AISDetected	This trap indicates that AIS Alarm is detected on the incoming signal.				

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description					
67	asxDS1AISCleared	This trap indicates that AIS Alarm is cleared on tincoming signal.					
68	asxDS1LOSDetected	This trap indicates that LOS Alarm is detected on the incoming signal.					
69	asxDS1LOSCleared	This trap indicates that LOS Alarm is cleared on the incoming signal.					
70	asxDS1LOFDetected	This trap indicates that LOF Alarm is detected on the incoming signal.					
71	asxDS1LOFCleared	This trap indicates that LOF Alarm is cleared on the incoming signal.					
74	asxDS3FERFDetected	This trap indicates that FERF Alarm is detected on the incoming signal.					
75	asxDS3FERFCleared	This trap indicates that FERF Alarm is cleared on the incoming signal.					
110	asxJ2YellowOn	This trap indicates that the specified J2 port is in the Yellow Alarm state. The Yellow alarm is declared if either LOF or LOS or AIS is detected and persists for 2.5+5 seconds.					
111	asxJ2YellowOff	This trap indicates that the specified J2 port is no longer in the Yellow Alarm state.					
112	asxJ2YellowDetected	This trap indicates that Yellow Alarm is detected on the incoming signal.					
113	asxJ2YellowCleared	This trap indicates that Yellow Alarm is cleared on the incoming signal.					
114	asxJ2AISDetected	This trap indicates that AIS Alarm is detected on the incoming signal.					
115	asxJ2AISCleared	This trap indicates that AIS Alarm is cleared on the incoming signal.					
116	asxJ2LOSDetected	This trap indicates that LOS Alarm is detected on the incoming signal.					
117	asxJ2LOSCleared	This trap indicates that LOS Alarm is cleared on the incoming signal.					

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
118	asxJ2LOFDetected	This trap indicates that LOF Alarm is detected on the incoming signal.
119	asxJ2LOFCleared	This trap indicates that LOF Alarm is cleared on the incoming signal.
1024	asxOutputQueueCongested	This trap indicates that the output queue for the given priority has exceeded its dedicated length, and has begun overflowing into the shared buffer space on the network module.
1025	asxOutputQueueCellLoss	This trap indicates that the output queue for the given priority has overflowed and cells have been dropped.
1026	as x Extended Mode Violation	This trap indicates that a series A or B network module was inserted into a switch board running in extended mode.
1027	asxNonextendedMode- Warning	This trap indicates that a series C or greater network module was inserted into a switch board running in non-extended mode.
1028	q2931AVRejectTrap	This trap is generated whenever any UNI3.x with AddressValidation enabled rejects a Setup Request call more than q2931AVRejectTrapThreshold times in any given q2931AVRejectTrapPeriod.
1029	crConfMemoryOflow	This trap is generated when the allocated call record memory (as indicated by crMemoryAllocated) is exceeded.
1030	crXfrPrimaryXfrFailed	This trap is generated when the call record transfer to the primary host (as indicated by crXfrPrimaryUrl) fails.
1031	crXfrSecondaryXfrFailed	This trap is generated when the call record transfer to the secondary host (as indicated by crXfrSecondaryUrl) fails.
1032	crConfMemAllocFail	This trap is generated when Callrecord functionality is unable to allocate memory as specified by crMemory-Allocated. This can happen when the crConfAdmin-Status changes state from "off" or when the switch reboots when Callrecords is configured "on".

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
1033	crGeneralFailure	This trap is generated when any of the callrecord related functionality fails for any reason. One example would be when the Callrecord Module fails to schedule an interval timer.
1034	asxDualScpSyncFailure	This trap indicates that automatic CDB synchronization is disabled due to failures.
1035	asxDualScpSwitchOver	This trap indicates that the backup SCP has taken control of the switch.
1036	asxDualScpHotSwap	This trap indicates that an SCP hot-swap insertion or removal has occurred.
1037	asxVPAISDetected	This trap indicates that the Alarm Indication Signal (AIS) is detected on the incoming (terminating) virtual path. This trap is generated once when the virtual path is declared to be in the active AIS state.
1038	asxVPAISCleared	This trap indicates that the Alarm Indication Signal (AIS) has been removed from the incoming (terminating) virtual path. This trap is generated once when the virtual path is declared to be in the inactive AIS state.
1039	asxVPRDIDetected	This trap indicates that the Remote Defect Indication (RDI) is detected on the incoming (terminating) virtual path. This trap is generated once when the virtual path is declared to be in the active RDI state.
1040	asxVPRDICleared	This trap indicates that the Remote Defect Indication (RDI) has been removed from the incoming (terminating) virtual path. This trap is generated once when the virtual path is declared to be in the inactive RDI state.
1041	asxNonextendedModeViolation	This trap indicates that a series D network module was inserted into a switch board running in non-extended mode. Multicast will not work on the series D module without removing all series B modules and rebooting the switch.

Table A.2 - SNMP Traps Supported on the MSC-900

Trap Number	Trap Name	Description
1042	asxUnsupportedNetwork- Module	This trap indicates that a unsupported network module was inserted into a switch.
1061	q2931CreationFailure	This trap is generated whenever a switch fails to create a UNI. This is most likely due to a resource limitation on the switch.
1062	sc Upl Netm Not Allowed	This trap indicates that an uplink netmod has been installed into an access netmod slot.
1063	sc Upl Net m Not All w Cleared	This trap indicates that a misplaced uplink netmod has been removed from an access netmod slot.
1068	asxPsCurrentDown	This trap alerts that one ATM switch power supply had a current failure. The power supply that failed is identified by the power supply index.
1069	asxPsCurrentUp	This trap alerts that one ATM switch power supply that had a current failure is now up. The power supply that is back up is identified by the power supply index.
1070	asxPs5VoltDown	This trap alerts that one ATM switch power supply had a +5V failure. The power supply that failed is identified by the power supply index.
1071	asxPs5VoltUp	This trap alerts that one ATM switch power supply that had a +5V failure is now up. The power supply that is back up is identified by the power supply index.

A.2.1 Adding SNMP Trap Destinations

To create one or more SNMP trap destinations on an MSC-900, log in to AMI and open a session on the MSC-900. Enter the following parameters:

```
configuration snmp trap new <ipaddress>
```

The <ipaddress> variable indicates the IP address of the SNMP trap destination that is to be created. Repeat this for as many SNMP trap destinations as needed. Traps are active as soon as they are set.

A.2.2 Displaying SNMP Trap Destinations

To list all of the current SNMP trap destinations, log in to AMI and open a session on the MSC-900. The SNMP traps supported by this MSC-900 are detailed in the FORE-Switch-MIB. Enter the following parameters:

configuration snmp trap show

Trap	Destination
1	192.88.243.18
2	198.29.16.14
3	198.29.16.18
4	198.29.23.39
5	198.29.31.130

A.2.3 Removing SNMP Trap Destinations

To delete one or more SNMP trap destinations for an MSC-900, log in to AMI and open a session on the MSC-900. Prior to deleting any trap that may need to be recreated later, as a precaution, a recommended practice is to list all trap destinations using AMI and either copy the screen or write down the destinations. To delete a trap, enter the following parameters:

```
configuration snmp trap delete <trap>
```

The <trap> variable indicates the SNMP trap number that is to be removed. Repeat this for as many SNMP trap destinations as needed.

APPENDIX B ForeThought PNNI

PNNI (Private Network Node Interface or Private Network-to-Network Interface) is a protocol defining interoperability between private ATM switches. PNNI defines both the routing and signalling standards for inter-switch interoperability.

This appendix provides a broad overview of FORE Systems' version of PNNI, ForeThought PNNI (FT-PNNI), and its use in a multiple-switch network. FT-PNNI is a scalable routing and signalling protocol used in networks containing multiple TNX switches. FT-PNNI simplifies large network topologies by organizing the nodes (switches) in that network into smaller groups.

It is this reorganization of the network topology that makes FT-PNNI's simplified routing possible. By segmenting a large network into smaller groups of nodes, FT-PNNI reduces the amount of topology information that those very nodes must maintain.

To understand FT-PNNI, it is important to first understand how a large network topology can be simplified by grouping the nodes contained therein. Once the organization of the physical network is determined, FT-PNNI's routing and signalling protocols can be better understood.

This appendix examines the routing and signalling that FT-PNNI utilizes. It also looks at the physical network and how it can be simplified through the use of peer groups.

B.1 FT-PNNI Routing

The FT-PNNI routing protocol serves to distribute topology and address reachability information between switches and groups of switches in a network. This topology and addressing information is used by switches to compute paths through the network. FT-PNNI routing has the following aspects:

- Hello Protocol
- Topology Database Exchange
- Flooding
- Hierarchical Routing

B.1.1 Hello Protocol

At regular intervals, each switch transmits a hello indication on each of its FT-PNNI routing channels. The time between these hello indications is called the Hello Indication Interval. When a switch receives a hello indication from one of its neighbors, it initializes the logical link (loglink) from that neighbor to itself in the topology database.

Since hello indications are sent periodically, loglinks are also refreshed periodically. Note that loglinks discovered as a result of hello indications are unidirectional. Each switch initializes unidirectional loglinks that have an adjacent switch as the destination.

B.1.2 Topology Database Exchange

At regular intervals, called the NSAP map indication interval, each switch sends to each of its neighbors a group of loglinks taken from its topology database. This is done to ensure that the topology databases of the switches stay synchronized.

B.1.3 Flooding

Whenever a new loglink is discovered by a switch, the switch immediately broadcasts the existence of this link to all of its neighbors. The neighboring switches then broadcast the existence of the same link to all of their neighbors. Very quickly, the existence of the new loglink is flooded immediately to all of the switches in the network.

Similarly, when a link goes down, that information is propagated through the network immediately. Also, when a *significant* change is seen in the metrics of a loglink, the latest state of the loglink is propagated immediately throughout the network.

B.1.4 Hierarchical Routing

Hierarchical routing is a method of routing that uses hierarchical addressing of nodes (switches), thereby enabling the logical partitioning of networks, topology information hiding between the partitions, and increased scalability of the routing protocol. FT-PNNI supports hierarchical node identifiers and, therefore, the hierarchical organization of networks.

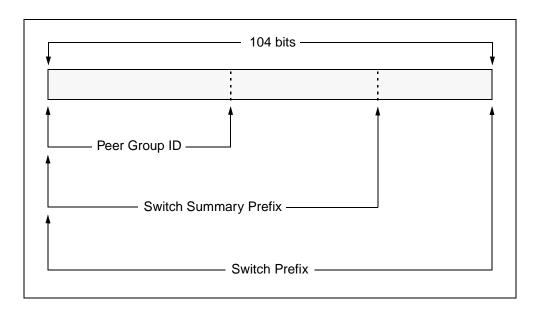


Figure B.1 - Example of a 13-byte Switch Prefix

B.1.4.1 Hierarchical Addressing

FT-PNNI uses private ATM address prefixes (NSAP prefixes) as node identifiers. FT-PNNI does not distinguish between node identifiers and reachability information. Thus, the IDs of nodes in the FT-PNNI addressing map are NSAP prefixes. In the default case at the lowest level peer group, the switches have a 13-byte prefix as their node ID and end systems (hosts) have a 19-byte prefix as their node ID.

B.1.4.1.1 Switch Prefix

Each switch in a FT-PNNI network is configured with a 13-byte prefix called the switch prefix. Hosts that are attached to the switch are presented with this prefix during ILMI address registration. In this way, end systems are configured with a private ATM address that includes the 13-byte switch prefix.

B.1.4.1.2 Switch Summary Prefix

Each switch is configured with a switch mask (*swmask*) which gives the length of the switch summary prefix within the switch prefix. The *swmask* gives the number of most significant bits of the switch prefix that constitute the switch summary prefix. Since all end system addresses attached to a switch have the same switch summary prefix, their reachability information can be summarized by this prefix (i.e., by the switch summary prefix).

B.1.4.1.3 Peer Group ID

Each switch is configured with a peer group mask (*pgmask*) which gives the length of the peer group ID within the switch summary prefix. The *pgmask* gives the number of most significant bits of the switch summary prefix that constitute the peer group ID.

The peer group ID is a unique identifier for every node (switch or end system) in a particular peer group. Every switch and host in peer group A, for example, will have the same peer group ID.

A simple example of summarizing by peer group ID can be seen in Figure B.3, where every switch (and end system, although not shown) in peer group A is identified starting with "A."

B.1.4.2 Path Computation

Path computation is performed on demand whenever FT-PNNI signalling requests a path to a given destination. The Bellman-Ford Shortest Path algorithm is used to compute the shortest path tree of all nodes in the topology with the local node as the source. The administrative weight metric in the loglinks is used as the minimizing criterion in computing the shortest path route. In the case of finding multiple equal-cost paths to a given destination, available cell rate is used to break ties.

B.2 The Physical Network

In an ATM network, data is sent and received over virtual circuits, or circuits that only exist when needed. This communication over these virtual circuits is made possible by signalling that occurs between the switches in the network.

In a network of TNX switches, any new addition to the topology is recognized immediately by all nodes (switches) having a direct connection to the new node. Then each switch that has recognized the new switch sends a message to each of its direct connections, and so on. Eventually, and within a very brief period of time, every switch in the network is aware of the new addition and the links by which that new addition can be reached. This topology is stored by each switch in a topology database.

In a small, local area network (LAN), the topology is relatively simple, meaning that the switches in the LAN have a relatively small topology database to maintain. As the LAN grows, however, and more switches are added, that same database can grow to be very large in a short period of time.

As this topology database grows, the amount of information a switch must consult when searching for an address also grows. In the end, this can result in delayed connection set-up, congestion in the network, and even lost data.

Figure B.2 depicts a typical ATM network, containing 21 TNX switches (()). The hierarchy of this network is flat, meaning that each switch must be aware of all the other switches in the network, as well as all the possible routes to those switches. As more switches are added to this network, the hierarchy will become more complex and the switches will have to contend with larger topology databases.

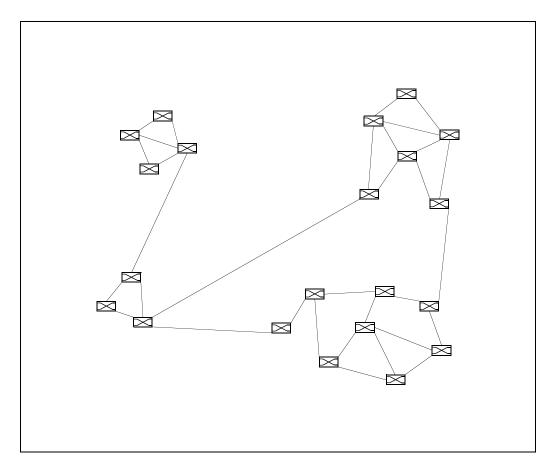


Figure B.2 - Private ATM Network with 21 Switches and 34 Bidirectional Links

It is in these large, single-level networks that FT-PNNI is most useful, because it lets you simplify large network topologies by creating a two-level hierarchy. In this hierarchy, communities of contiguous switches are grouped together and they are collectively summarized by a single, logical node.

Figure B.3 shows the same network as in Figure B.2 after being organized into peer groups, now having a two-level hierarchy. The subsections that follow explain the organization of these peer groups, how they simplify the overall network topology, and how they change the logical view of the network.

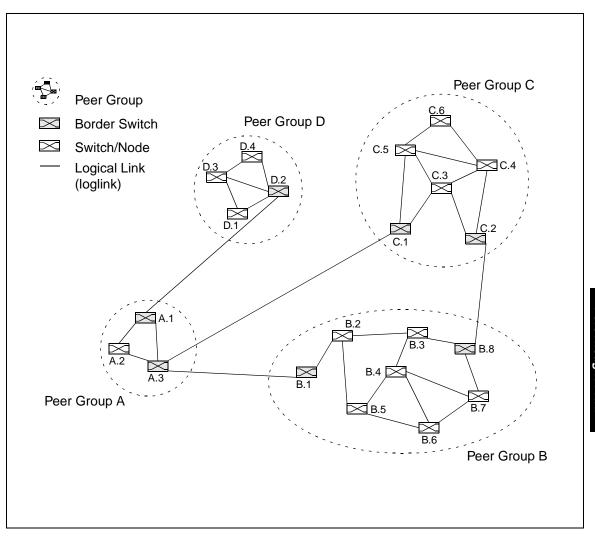


Figure B.3 - Example of FT-PNNI Hierarchy Showing Lowest-Level Peer Groups

B.2.1 Peer Groups

The FT-PNNI hierarchy begins with a network of switches, organized into peer groups. A peer group is a collection of interconnected switches that are organized into a group. Peer group organization can be determined by a network administrator, but switches that are located close to one another are usually made into a peer group.

The network shown in Figure B.3 is organized into four lowest-level peer groups: A, B, C, and D. The switches within a certain peer group are numbered according to that particular group. For example, the switches in peer group A are identified as A.1, A.2, and A.3.

Peer groups have a peer group identifier (ID), assigned at configuration time and exchanged in hello indication messages. Switches can determine in which peer group they are located by comparing these peer group IDs.

Switches in a peer group are aware of the topology of their own peer group and the existence of all other peer groups. They recognize the links between their peer group and others, but they are not aware of internal topology information of other peer groups. Instead, the topologies of other peer groups are summarized as a single, reachable location, known as a peer group summary node (PGSN).

B.2.2 Peer Group Topology

To maintain an accurate and updated view of its relative location and status, a switch sends hello indication messages to every other switch with which it has a direct connection. These hello indications contain the switch prefix, peer group membership information, and link metrics (attributes) for the physical link between the two switches.

Through this exchange of messages, each switch learns which switches are its immediate neighbors, to what peer groups they belong, and whether or not the link between itself and its neighbors is valid.

B.2.3 Border Switches

A border switch is any switch that has at least one link to a switch in another peer group. Border switches play an important role in FT-PNNI because they are responsible for summarizing reachability information for their respective peer groups, appropriately filtering the flow of topology database information across peer group boundaries, and building the lowest level source route for call setups entering the peer group.

B.2.4 Peer Group Summary Node (PGSN)

A PGSN is a hypothetical or imaginary node that summarizes a peer group's reachability information. The PGSN has the peer group ID of its peer group as its switch summary prefix. Each border switch in the peer group advertises a logical link (loglink) to the PGSN. The PGSN is a logical representation of the switches contained in a peer group.

B.2.5 Backbone Topology

Loglinks between border switches and loglinks from border switches to PGSNs are called backbone links and considered part of the backbone topology. Information regarding these backbone links is propagated across peer group boundaries during database exchange and flooding.

B.2.6 Single Switch Perspective

The main reason for grouping switches in large networks is to simplify each individual switch's view of the topology. For example, each switch in peer group A is aware of every other switch in peer group A, the border switches in the rest of the network, the links between them, and the backbone topology. Switches in peer group A are not aware, however, of the internal topology of other peer groups. Instead, individual switches see a PGSN (see Figure B.4).

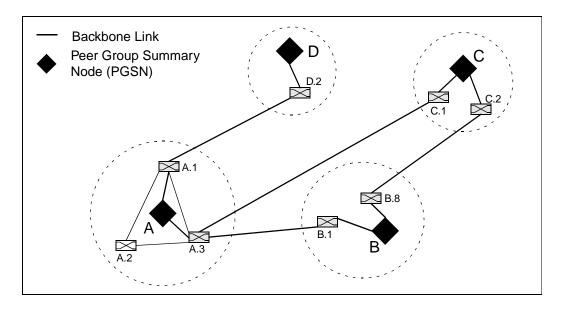


Figure B.4 - View of the Network from Switches in Peer Group A

APPENDIX C UNI Signalling

This appendix contains information that pertains to the UNI signalling portion of *ForeThought*-SP 1.0.x software. This UNI signalling information is described in the following sections:

- Section C.1 VCI Allocation Range
- Section C.2 Non-path Associated Signalling Enhancements
- Section C.3 UNI Auto Configuration Procedures
- Section C.4 Allowable Combination of Traffic Parameters

C.1 VCI Allocation Range

The VCI allocation range is that range of VCI values in a path within which a UNI signalling channel will allocate VCs. This range defaults to the VCI-range of the path containing the signalling channel associated with the UNI.



If the UNI is a link-scope UNI, this range applies to the containing path and all other paths within which the UNI has allocated connections.

A VCI allocation range can be specified when a UNI is created using the ATM Management Interface (AMI) command conf uni new. The range is computed differently depending on whether or not you have ILMI running.

C.1.1 Determining the VCI Allocation Range with ILMI Down

If ILMI is not operational and if the VCI allocation range of the UNI is specified, then the range is determined as an intersection of the following:

- the VCI-range of the path containing the signalling channel associated with the UNI
- the range specified for the UNI

For example, if you create a UNI with ILMI down on a VP that has no VCIs reserved, and you specify a VCI range:

```
conf uni new 1A3 65 -ilmi down -minvci 70 -maxvci 200
```

the actual range is computed as an intersection based on the range you enter (70 - 200) and the range available in that VP (1 - 511).

To display what the actual VCI range is, enter the following command:

conf uni show advanced

					Admin	Admin	Oper	Oper		
Port	VPI	SigVCI	ILMIVCI	SigAAL	MinVCI	MaxVCI	MinVCI	MaxVCI	OrigCost	TermCost
1A1	0	5	16	5	32	511	32	511	100	100
1A2	0	5	16	5	32	511	[32]	[511]	100	0
1A3	0	5	16	5	32	511	[32]	[511]	100	0
1A3	65	5	16	5	70	200	70	200	100	0
1A4	0	5	16	5	32	511	[32]	[511]	100	0
1CTL	0	5	16	5	32	1023	32	1023	100	0

In this display, the range that you have requested is shown in the Admin MinVCI and Admin MaxVCI fields. The range that is computed and allowed is displayed in the Oper MinVCI and Oper MinVCI fields. In this example, no VCIs are reserved on VP 65, so the requested range of 70 - 200 is allowed.

In another example, if there are VCIs already reserved on VP 65, as shown below in the MinVCI and MaxVCI fields:

<pre>localhost::configuration vpt></pre>			show						
Input Output									
Port	VPI	Port VPI ResBW		CurBW	MinVCI	MaxVCI	VCs	Protocol	
1A1	0	termi	nate	N/A	37.3K	1	511	35	pvc
1A2	0	termi	nate	N/A	0.8K	1	511	6	pvc
1A3	0	termi	nate	N/A	0.8K	1	511	6	pvc
1A3	65	terminate		N/A	0.8K	1	120	6	pvc
1A4	0	termi	nate	N/A	0.8K	1	511	6	pvc
1CTL	0	termi	nate	N/A	45.4K	1	1023	85	pvc
origi	nate	1A1	0	N/A	1.7K	1	511	7	pvc
origi	nate	1A2	0	N/A	0.8K	1	511	6	pvc
origi	nate	1A3	0	N/A	0.8K	1	511	6	pvc
origi	nate	1A3	65	N/A	0.8K	1	120	6	pvc
originate		1A4	0	N/A	0.8K	1	511	6	pvc
Press return for more, q to quit: q									

and you enter the following:

conf uni new 1A3 65 -ilmi down -minvci 70 -maxvci 200

then the actual range is computed as an intersection based on the range you enter (70 - 200) and the range available in that VP (1 - 120).

To see what the actual VCI range is, enter the following command:

conf uni show advanced

					Admin	Admin	Oper	Oper		
Port	VPI	SigVCI	ILMIVCI	SigAAL	MinVCI	MaxVCI	MinVCI	MaxVCI	OrigCost	TermCost
1A1	0	5	16	5	32	511	32	511	100	100
1A2	0	5	16	5	32	511	[32]	[511]	100	0
1A3	0	5	16	5	32	511	[32]	[511]	100	0
1A3	65	5	16	5	70	200	70	120	100	0
1A4	0	5	16	5	32	511	[32]	[511]	100	0
1CTL	0	5	16	5	32	1023	32	1023	100	0

In this example, since VCIs 1-120 are already reserved on VP 65, the range of 70-120 is the actual range allowed.

C.1.2 Determining the VCI Allocation Range with ILMI Up

If ILMI is operational, then the peer's VCI allocation range is obtained by reading its ILMI MIB variable atmfAtmLayerMaxVccs. The VCI allocation range of the UNI is then determined as an intersection of the following:

- the VCI-range of the path containing the signalling channel associated with the UNI
- the range specified for the UNI (if applicable)
- the VCI-range of the peer signalling entity

For example, if your peer has no VCIs reserved, and you create a UNI with ILMI up on a VP that has no VCIs reserved, and you specify a VCI range:

```
conf uni new 1A3 65 -ilmi up -minvci 70 -maxvci 200
```

the actual range is computed as an intersection based on the range you enter (70 - 200), the range available in that VP (1 - 511), and the range of the peer signalling entity (32 - 511).

To display what the actual VCI range is, enter the following command:

conf uni show advanced

					Admin	Admin	Oper	Oper		
Port	VPI	SigVCI	ILMIVCI	SigAAL	MinVCI	MaxVCI	MinVCI	MaxVCI	OrigCost	TermCost
1A1	0	5	16	5	32	511	32	511	100	100
1A2	0	5	16	5	32	511	[32]	[511]	100	0
1A3	0	5	16	5	32	511	[32]	[511]	100	0
1A3	65	5	16	5	70	200	70	200	100	0
1A4	0	5	16	5	32	511	[32]	[511]	100	0
1CTL	0	5	16	5	32	1023	32	1023	100	0

In this display, the range that you have requested is shown in the Admin MinVCI and Admin MaxVCI fields. The range that is computed and allowed is displayed in the Oper MinVCI and Oper MinVCI fields. In this example, no VCIs are reserved on VP 65 and the signalling peer supports the range of 32 - 511, so the requested range of 70 - 200 is allowed.

In another example, if the peer supports the range of 32 - 511, but there are VCIs already reserved on VP 65, as shown below in the MinVCI and MaxVCI fields:

localh	ost::	config	uratio	on vpt>	show				
Input		Outpu	.t						
Port	VPI	Port	VPI	ResBW	CurBW	MinVCI	MaxVCI	VCs	Protocol
1A1	0	termi	nate	N/A	37.3K	1	511	35	pvc
1A2	0	termi	nate	N/A	0.8K	1	511	6	pvc
1A3	0	termi	nate	N/A	0.8K	1	511	6	pvc
1A3	65	termi	nate	N/A	0.8K	1	120	6	pvc
1A4	0	termi	nate	N/A	0.8K	1	511	6	pvc
1CTL	0	termi	nate	N/A	45.4K	1	1023	85	pvc
origi	nate	1A1	0	N/A	1.7K	1	511	7	pvc
originate		1A2	0	N/A	0.8K	1	511	6	pvc
originate		1A3	0	N/A	0.8K	1	511	6	pvc
originate		1A3	65	N/A	0.8K	1	120	6	pvc
originate		1A4	0	N/A	0.8K	1	511	6	pvc
Press	retur	n for	more,	q to qu	uit: q				

and you enter the following:

conf uni new 1A3 65 -ilmi up -minvci 70 -maxvci 200

then the actual range is computed as an intersection based on the range you enter (70 - 200), the range available in that VP (1 - 120), and the range supported by the peer (32 - 511).

To see what the actual VCI range is, enter the following command:

conf uni show advanced

					Admin	Admin	Oper	Oper		
Port	VPI	SigVCI	ILMIVCI	SigAAL	MinVCI	MaxVCI	MinVCI	MaxVCI	OrigCost	TermCost
1A1	0	5	16	5	32	511	32	511	100	100
1A2	0	5	16	5	32	511	[32]	[511]	100	0
1A3	0	5	16	5	32	511	[32]	[511]	100	0
1A3	65	5	16	5	70	200	70	120	100	0
1A4	0	5	16	5	32	511	[32]	[511]	100	0
1CTL	0	5	16	5	32	1023	32	1023	100	0

In this example, since the peer supports VCIs 32 - 511, but VCIs 1 -120 are already reserved on VP 65, the range of 70 - 120 is the actual range allowed.

In a third example, if the peer only supports the range of 32 - 180, and VCIs 1 -120 are already reserved on VP 65, and you enter the following:

then the actual range is computed as an intersection based on the range you enter (70 - 200), the range available in that VP (1 - 120), and the range supported by the peer (32 - 180).

To see what the actual VCI range is, enter the following command:

conf uni show advanced

					Admin	Admin	Oper	Oper		
Port	VPI	SigVCI	ILMIVCI	SigAAL	MinVCI	MaxVCI	MinVCI	MaxVCI	OrigCost	TermCost
1A1	0	5	16	5	32	511	32	511	100	100
1A2	0	5	16	5	32	511	[32]	[511]	100	0
1A3	0	5	16	5	32	511	[32]	[511]	100	0
1A3	65	5	16	5	70	200	70	120	100	0
1A4	0	5	16	5	32	511	[32]	[511]	100	0
1CTL	0	5	16	5	32	1023	32	1023	100	0

In this example, since the peer supports VCIs 32 - 180, and VCIs 1 -120 are already reserved on VP 65, the range of 70 - 120 is the actual range allowed.

C.2 Non-path Associated Signalling Enhancements

In *ForeThought*-SP 1.0.x, a signalling channel can control VCs within other VPs on the same physical link.

A UNI can be configured to control either VCs within its own VP (VP-scope, the default) or to control VCs in its own VP, in addition to controlling VCs in other VPs on the link (link-scope).

The following criteria must be satisfied in order to use link-scope UNIs:

- There can be only one link-scope UNI on a given link.
- A link-scope UNI must be created within VP 0 on the given link.
- The MSC-900 must be in extended mode.
- The path containing the link-scope UNI must be an elastic path.
- A backplane UNI cannot be a link-scope UNI.

C.2.1 VC-Space

The VC-space of a UNI is the VPI/VCI space within which a UNI can allocate connections. This is determined by the allocation scope of the UNI and the VCI allocation range (see Section C.1) for the UNI.

The VC-space of a VP-scope UNI is the VCI allocation range of the UNI applied only to the path that contains the signalling channel associated with the UNI.

The VC-space of a link-scope UNI is comprised of the aggregate VCI allocation range of the UNI applied to all paths controlled by the UNI (i.e., the path containing the signalling channel and all other paths in which the UNI has allocated connections).

C.2.2 Dynamic Paths

A link-scope UNI allocates new virtual channel connections based on the VC space availability of the link-scope UNI at that time. A link-scope UNI allocates connections in a new originating/terminating path (called a "dynamic path," because signalling uses such a path on the fly) when its current VC-space is exhausted.



Every time a a link-scope UNI uses a new path for allocating connections, its VC-space increases.

To avoid conflicts of ownership of the path, a link-scope UNI never uses an existing provisioned terminating/originating path (apart from the path that contains the signalling channel) to allocate virtual channel connections.

Dynamic paths inherit the characteristics of the path that contains the signalling channel (called the "containing path"). Therefore, dynamic paths have the same VCI-range as the containing path and are elastic paths.

To see which paths are dynamic paths, look at the Protocol field associated with each path under the AMI command conf vpt show. This field displays q2931 for dynamic paths and displays pvc for provisioned terminating/originating paths.

A dynamic path is controlled only by the link-scope UNI that uses it. Therefore, a dynamic path is not open to any provisioning such as the creation of a PVC connection, SPANS signalling path, etc.

C.3 UNI Auto Configuration Procedures

UNI 3.1 signalling support is provided in *ForeThought-SP* 1.0.x in addition to UNI 3.0 signalling support. The MSC-900 can be configured to detect which connections are UNI 3.0 and which are UNI 3.1 automatically. The MSC-900 also translates between connections of different types.

When a UNI signalling channel is created using conf uni new uni30 or conf uni new uni31, then the corresponding SSCOP stack is initialized and the specified version is used for that UNI. At the same time, a get request is sent for the atmfAtmLayerUniVersion MIB variable. If the peer does not support this MIB variable and returns a no such name response, it is assumed that the peer is using UNI 3.0. If a version mismatch occurs, an error message is displayed on the syslog. This request is repeated until a response comes from the peer or until the ILMI keep alive request times out. In any case, no change is made to the UNI that was explicitly configured.

When a UNI signalling channel is created using conf uni new auto, then the UNI creation is delayed until a response is obtained from the peer or until the ILMI keep alive requests time out. Based on the peer response, the appropriate SSCOP stack is initialized. If the peer does not support the atmfAtmLayerUniVersion variable, then the MSC-900 software defaults to UNI 3.0. Similarly, if the ILMI keep alive requests time out, then the MSC-900 software defaults to UNI 3.0.

The following two tables summarize the above information. Table A.1 highlights the actions taken based on the configurations of the FORE devices on both sides of an interface.

Table C.1 - Action Taken Based on Both Devices' UNI Configurations

Configured UNI Version	Peer's Configured UNI Version	Action Taken
UNI 3.0	UNI 3.0	SSCOP 30 stack initialized
	UNI 3.1	SSCOP 30 stack initialized; error message displayed on syslog
	Auto	SSCOP 30 stack initialized; peer changes version to UNI 3.0
UNI 3.1	UNI 3.0	SSCOP 31 stack initialized
	UNI 3.1	SSCOP 31 stack initialized; error message displayed on syslog
	Auto	SSCOP 31 stack initialized; peer changes version to UNI 3.1
Auto	UNI 3.0	SSCOP 30 stack initialized
	UNI 3.1	SSCOP 31 stack initialized
	Auto	Both stacks initialized as SSCOP 31

An alternative way of looking at the auto configuration procedure is to look at the atmfAtmLayerUniVersion variable. Table A.2 shows actions taken based on the atmfAtmLayerUniVersion variable supported by the peer device.

Table C.2 - Action Taken Based on the Peer's Supported MIB Variable

Configured UNI Version	atmfAtmLayer UNI Version	Action Taken
UNI 3.0	Version 3.0	SSCOP 30 stack initialized
	Version3.1	SSCOP 30 stack initialized; error message displayed on syslog
	Version 2.0	SSCOP 30 stack initialized
	MIB variable not supported	SSCOP 30 stack initialized
UNI 3.1	Version 3.0	SSCOP 31 stack initialized; error message displayed on syslog
	Version3.1	SSCOP 31 stack initialized
	Version 2.0	SSCOP 30 stack initialized
	MIB variable not supported	SSCOP 31 stack initialized; error message displayed on syslog
Auto	Version 3.0	SSCOP 30 stack initialized
	Version3.1	SSCOP 31 stack initialized
	Version 2.0	SSCOP 30 stack initialized
	MIB variable not supported	SSCOP 30 stack initialized

C.4 Allowable Combination of Traffic Parameters

FORE Systems supports a subset of traffic contracts which are specified in Table F.1 of the UNI 3.1 specification. These traffic contracts are enforced on both signalled connections as well on PNNI SPVCs. The following are the allowable contracts supported by the MSC-900.

• CBR: PCR (0), PCR (0+1)

• CBR: PCR (0+1)

• VBR: PCR (0+1), SCR, MBS (0)

• VBR: PCR (0+1),SCR, MBS (0+1)

• UBR: PCR (0+1)



Traffic contracts that do not completely characterize the traffic characteristics of a given service category are incomplete and not supported. (For example, a VBR contract that is received with an ATM traffic descriptor information element that does not specify SCR and MBS cell rates does not fully characterize a VBR contract. Such a contract is, therefore, rejected by the MSC-900).

Table C.3 shows the subset of the UNI 3.1 allowable combination of traffic parameters supported by *ForeThought-SP* 1.0.x:

Table C.3 - UNI 3.1 Allowable Combination of Traffic Parameters in *ForeThought-SP 1.0.x*

Broadband Bearer Class	A	X	С	X	С	X	A	X	С	X
Traffic Type		CBR		&		&		CBR		&
Timing Required		Y		&&		&&		Y		&&
PCR (CLP=0)	S	S								
PCR (CLP=0+1)	S	S	S	S	S	S	S	S	S	S
SCR and MBS (CLP=0)			S	S						
SCR and MBS (CLP=0+1)					S	S				
Best Effort									S	S
Tagging	Y/N	Y/N	Y/N	Y/N	N	N	N	N	N	N
QoS Class	*	*	*	*	*	*	*	*	0	0

The following are used in the table above:

- Y Yes
- N No
- **S** Specified
- Y/N Either Yes or No is allowed.
- * All QoS classes are supported.
 - & The parameter is coded to either "no indication," or "VBR," or, for UNI 3.1, Octet 5a (Traffic Type / Timing Required) is absent; these 3 codings are treated as equivalent.
- **&&** The parameter is coded to either "no indication," or "No," or, for UNI 3.1, Octet 5a (Traffic Type / Timing Required) is absent; these 3 codings are treated as equivalent.

A blank entry indicates that the parameter is not present.

UNI Signalling

Acronyms

The networking terms in the following list are defined in the Glossary of this manual. Glossary items are listed alphabetically according to the full term.

AAL ATM Adaptation Layer
ABR Available Bit Rate

ACM Address Complete Message

ACR Allowable Cell Rate

ADPCM Adaptive Differential Pulse Code Modulation

AHFG ATM-attached Host Functional Group

AIMUX ATM Inverse Multiplexing
AIS Alarm Indication Signal
AMI Alternate Mark Inversion
AMI ATM Management Interface

ANSI American National Standards Institute
APCM Adaptive Pulse Code Modulation
API Application Program Interface

APP Application Program

APS Automatic Protection Switching
ARP Address Resolution Protocol

ASCII American Standard Code for Information Interchange

ATDM Asynchronous Time Division Multiplexing

ATM Asynchronous Transfer Mode
AUI Attachment User Interface
BBZS Bipolar 8 Zero Substitution

BCOB Broadband Connection Oriented Bearer

BCOB-A Bearer Class A
BCOB-C Bearer Class C
BCOB-X Bearer Class X

BECN Backward Explicit Congestion Notification

BER Bit Error Rate

BES Bursty Errored SecondsBGP Border Gateway ProtocolB-ISDN Inter-Carrier Interface.

BIP Bit Interleaved Parity

B-ISDN Broadband Integrated Services Digital Network

B-ISUP Broadband ISDN User's Part

Acronyms

BITS Building Integrated Timing Supply

BNC Bayonet-Neill-Concelman
BPDU Bridge Protocol Data Unit

bps Bits per SecondBPV Bipolar Violation

B-TE Broadband Terminal Equipment
BUS Broadcast and Unknown Server
CAC Connection Admission Control
CAS Channel Associated Signaling

CBDS Connectionless Broadband Data Service

CBR Constant Bit Rate

CCITT International Telephone and Telegraph Consultative Committee

CCS Common Channel Signaling

CDV Cell Delay Variation
CE Connection Endpoint

CEI Connection Endpoint Identifier
CES Circuit Emulation Service

CGA Carrier Group Alarm
CIP Carrier Identification Parameter

CIR Committed Information Rate

CLIP Classical IP
CLP Cell Loss Priority
CLR Cell Loss Ratio-1-15
CLS Connectionless service

CMIP Common Management Interface Protocol

CMR Cell Misinsertion Rate

CPE Customer Premise Equipment

CRA Cell Rate Adaptation
CRC Cyclic Redundancy Check

CRS Cell Relay Service
CS Controlled Slip, or

CSU Channel Service Unit
CTD Cell Transfer Delay

CTS Clear To Send

DACS Digital Access and Cross-Connect System
DARPA Defense Advanced Research Projects Agency

DCC Data Country Code

DCS Data Communications Equipment
DCS Digital Cross-connect System
DES Destination End Station

DFA DXI Frame Address

DLCI Data Link Connection Identifier

DNS Domain Naming System

DSn Digital Standard n (n=0, 1, 1C, 2, and 3)

DSR Data Set Ready

DTE Data Terminal Equipment
DTR Data Terminal Ready

EEPROM Electrically Erasable Programmable Read Only Memory

EFCI Explicit Forward Congestion Indication

EGP Exterior Gateway Protocol

EIA Electronics Industries Association

EISA Extended Industry Standard Architecture

EMI Emulated Local Area Network Electromagnetic Interference

EPD Early Packet Discard

EPROM Erasable Programmable Read Only Memory

EQL Equalization
ER Explicit Rate
ES End System, or

Errored Second

ESF Extended Super Frame ESI End System Identifier

EXZ Excessive Zeroes (Error Event)

FC Face Contact

FCC Federal Communications Commission

FCS Frame Check Sequence

FDDI Fiber Distributed Data Interface
FDM Frequency Division Multiplexing

FEBE Far End Block Error
FEC Forward Error Correction

FECN Forward Explicit Congestion Notification

FERF Far End Receive Failure
FIFO First-In, First-Out
FRS Frame-Relay Service
FTP File Transfer Protocol
FT-PNNI ForeThought PNNI
FUNI Frame-Based UNI

GCAC Generic Connection Admission Control

GCRA Generic Cell Rate Algorithm

GFC Generic Flow Control HDB3 High Density Bipolar

HDLC High Level Data Link Control

HEC Header Error Control

HIPPI High Performance Parallel Interface

HSSI High-Speed Serial Interface

Acronyms

ICMP Internet Control Message Protocol

IDU Interface Data Unit

IEEE Institute of Electrical and Electronics Engineers

IETF Internet Engineering Task Force
ILMI Interim Local Management Interface

IP Internet Protocol

IPX Internetwork Packet Exchange

IS Intermediate system

ISDN Integrated Services Digital Network
ISO International Standards Organization

ITU-T International Telecommunication Union Telecommunication

IWF Interworking Function IXC Interexchange Carriers

JPEG Joint Photographic Experts Group

Kbps Kilobits per second
LAN Local Area Network
LAN Emulation

LAPB Link Access Procedure, Balanced Local Access and Transport Area

LEO Line Build Out
LCV Line Code Violations

LE ARP LAN Emulation Address Resolution Protocol

LEC LAN Emulation Client

LECS LAN Emulation Configuration Server

LES LAN Emulation Server
LLC Logical Link Control
LOF Loss Of Frame
LOP Loss Of Pointer

LOS Loss Of Signal
LSB Least Significant Bit
MAC Media Access Control
MAN Metropolitan Area Network
MAU Media Attachment Unit
MBS Maximum Burst Size

MCDV Maximum Cell Delay Variance
MCLR Maximum Cell Loss Ratio

MCR Minimum Cell Rate

MCTDMaximum Cell Transfer DelayMIBManagement Information BaseMICMedia Interface Connector

MID Message Identifier

MMF Multimode Fiber Optic Cable
MPEG Motion Picture Experts Group

MPOA Multiprotocol over ATM
MSB Most Significant Bit

MTU Maximum Transmission Unit
NM Network Management Entity
NML Network Management Layer
NMS Network Management Station

NNI Network-to-Network Interface or Network Node Interface

NPC Network Parameter Control

NRZ Non Return to Zero

NRZI Non Return to Zero Inverted
NSAP Network Service Access Point
NTSC National TV Standards Committee
OAM Operation and Maintenance Cell

OC-n Optical Carrier level-n
OID Object Identifier
OOF Out-of-Frame

OSI Open Systems Interconnection
OSPF Open Shortest Path First Protocol
OUI Organizationally Unique Identifier
PAD Packet Assembler Disassembler

PAL Phase Alternate Line
PBX Private Branch Exchange

PCI Peripheral Component Interconnect

PCM Pulse Code Modulation

PCR Peak Cell Rate

PDN Public Data Network
PDU Protocol Data Unit
PHY Physical Layer

ping Packet Internet Groper

PLCP Physical Layer Convergence Protocol

PLP Packet Level Protocol
PM Physical Medium

PMD Physical Medium Dependent

PNNI Private Network Node Interface or Private Network-to-Network Interface

PPD Partial Packet Discard
PPP Point-to-Point Protocol

PROM Programmable Read-Only Memory

PRS Primary Reference Source
PSN Packet Switched Network

PT Payload Type

PVC Permanent Virtual Circuit (or Channel)
PVCC Permanent Virtual Channel Connection
PVPC Permanent Virtual Path Connection

Acronyms

QD Queuing Delay
QoS Quality of Service
RD Routing Domain
RFCs Requests For Comment
RFI Radio Frequency Interference
RIP Routing Information Protocol
RISC Reduced Instruction Set Computer

RTS Request To Send
SA Source Address
SA Source MAC Address
SAP Service Access Point

SAR Segmentation And Reassembly

SC Structured Cabling, or

Structured Connectors, or

Stick and Click

SCR Sustainable Cell Rate

SCSI Small Computer Systems Interface
SDLC Synchronous Data Link Control

SDU Service Data Unit

SEAL Simple and Efficient Adaptation Layer
SECAM Systeme En Coleur Avec Memoire

SEL Selector

SES Severely Errored Seconds

SF Super Frame

SGMP Simple Gateway Management Protocol

SIR Sustained Information Rate

SLIP Serial Line IP

SMDS Switched Multimegabit Data Service

SMF Single Mode Fiber

SMTP Simple Mail Transfer Protocol
SNA Systems Network Architecture
SNAP SubNetwork Access Protocol
SNI Subscriber Network Interface

SNMP Simple Network Management Protocol

SONET Synchronous Optical Network

SPANS Simple Protocol for ATM Network Signalling

SPARC Scalable Processor Architecture Reduced instruction set Computer

SPE Synchronous Payload Envelope

SPVC Smart PVC

SS7 Signaling System No. 7

SSCOP Service Specific Connection Oriented Protocol

SSCS Service Specific Convergence Sublayer

Straight Tip, or

Stick and Turn

STM Synchronous Transfer Mode

STP Shielded Twisted Pair, Spanning Tree Protocol

STS Synchronous Transport Signal

SVC Switched Virtual Circuit (or Channel)
SVCC Switched Virtual Channel Connection
SVPC Switched Virtual Path Connection

TAXI Transparent Asynchronous Transmitter/Receiver Interface

TC Transmission Convergence
TCP Transmission Control Protocol

TCP/IP Transmission Control Protocol/Internet Protocol

TCR Tagged Cell Rate

TCS Transmission Convergence Sublayer

TDM Time Division Multiplexing

TE Terminal Equipment

TFTP Trivial File Transfer Protocol

TM Traffic Management
UAS Unavailable Seconds
UBR Unspecified Bit Rate
UDP User Datagram Protocol
UNI User-to-Network Interface
UPC Usage Parameter Control

UTOPIA Universal Test & Operations Interface for ATM

UTP Unshielded Twisted Pair

VBR Variable Bit Rate

VC Virtual Channel (or Circuit)
VCC Virtual Channel Connection
VCI Virtual Channel Identifier
VCL Virtual Channel Link
VINES Virtual Network Software
VLAN Virtual Local Area Network

VP Virtual Path

VPC Virtual Path Connection
VPDN Virtual Private Data Network

VPI Virtual Path Identifier
VPL Virtual Path Link
VPN Virtual Private Network
VPT Virtual Path Terminator

VS/VD Virtual Source/Virtual Destination

VT Virtual Tributary
WAN Wide-Area Network

ZBTSI Zero Byte Time Slot Interchange

Acronyms

Glossary

10Base-T - a 10 Mbps baseband Ethernet specification utilizing twisted-pair cabling (Category 3, 4, or 5). 10BaseT, which is part of the IEEE 802.3 specification, has a distance limit of approximately 100 meters per segment.

802.1d Spanning Tree Bridging - the IEEE standard for bridging; a MAC layer standard for transparently connecting two or more LANs (often called subnetworks) that are running the same protocols and cabling. This arrangement creates an extended network, in which any two workstations on the linked LANs can share data.

802.3 Ethernet - the IEEE standard for Ethernet; a physical-layer standard that uses the CSMA/CD access method on a bus-topology LAN.

802.5 Token Ring - the IEEE physical-layer standard that uses the token-passing access method on a ring-topology LAN.

AAL Connection - an association established by the AAL between two or more next higher layer entities.

Adapter - A fitting that supplies a passage between two sets of equipment when they cannot be directly interconnected.

Adaptive Differential Pulse Code Modulation (ADPCM) - A technique that allows analog voice signals to be carried on a 32K bps digital channel. Sampling is done at 8Hz with 4 bits used to describe the difference between adjacent samples.

Adaptive Pulse Code Modulation (APCM) - A technique that effectively reduces occupied bandwidth per active speaker by reducing sampling rates during periods of overflow peak traffic.

Address - A unique identity of each network station on a LAN or WAN.

 $\label{lem:Address Complete Message (ACM) - A B-ISUP call control message from the receiving exchange to sending exchange indicating the completion of address information.}$

Address Mask - a bit mask used to identify which bits in an address (usually an IP address) are network significant, subnet significant, and host significant portions of the complete address. This mask is also known as the subnet mask because the subnetwork portion of the address can be determined by comparing the binary version of the mask to an IP address in that subnet. The mask holds the same number of bits as the protocol address it references.

Address Prefix - A string of 0 or more bits up to a maximum of 152 bits that is the lead portion of one or more ATM addresses.

Address Resolution - The procedure by which a client associates a LAN destination with the ATM address of another client or the BUS.

Address Resolution Protocol (ARP) - a method used to resolve higher level protocol addressing (such as IP) into the appropriate header data required for ATM; i.e., port, VPI, and VCI; also defines the AAL type to be used.

Agent - a component of network- and desktop-management software, such as SNMP, that gathers information from MIBs.

alarm - an unsolicited message from a device, typically indicating a problem with the system that requires attention.

Alarm Indication Signal (AIS) - In T1, an all ones condition used to alert a receiver that its incoming signal (or frame) has been lost. The loss of signal or frame is detected at the receiving end, and the failed signal is replaced by all the ones condition which the receiver interprets as an AIS. The normal response to this is AIS is for the receiving end to generate a yellow alarm signal as part of its transmission towards the faulty end. (The AIS itself is sometimes called a Blue Signal).

A-Law - The PCM coding and companding standard used in Europe.

Allowable Cell Rate (ACR) - parameter defined by the ATM Forum for ATM traffic management. ACR varies between the MCR and the PCR, and is dynamically controlled using congestion control mechanisms.

Alternate Mark Inversion (AMI) - A line coding format used on T1 facilities that transmits ones by alternate positive and negative pulses.

Alternate Routing - A mechanism that supports the use of a new path after an attempt to set up a connection along a previously selected path fails.

American National Standards Institute (ANSI) - a private organization that coordinates the setting and approval of some U.S. standards. It also represents the United States to the International Standards Organization.

American Standard Code for Information Interchange (ASCII) - a standard character set that (typically) assigns a 7-bit sequence to each letter, number, and selected control characters.

AppleTalk - a networking protocol developed by Apple Computer for communication between Apple's products and other computers. Independent of the network layer, AppleTalk runs on LocalTalk, EtherTalk and TokenTalk.

Application Layer - Layer seven of the ISO reference model; provides the end-user interface.

Application Program (APP) - a complete, self-contained program that performs a specific function directly for the user.

Application Program Interface (API) - a language format that defines how a program can be made to interact with another program, service, or other software; it allows users to develop custom interfaces with FORE products.

Assigned Cell - a cell that provides a service to an upper layer entity or ATM Layer Management entity (ATMM-entity).

 \mbox{asxmon} - a FORE program that repeatedly displays the state of the switch and its active ports.

Asynchronous Time Division Multiplexing (ATDM) - a multiplexing technique in which a transmission capability is organized into a priori, unassigned time slots. The time slots are assigned to cells upon request of each application's instantaneous real need.

Asynchronous Transfer Mode (ATM) - a transfer mode in which the information is organized into cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.

ATM Adaptation Layer (AAL) - the AAL divides user information into segments suitable for packaging into a series of ATM cells. AAL layer types are used as follows:

AAL-1 - constant bit rate, time-dependent traffic such as voice and video

AAL-2 - still undefined; a placeholder for variable bit rate video transmission

AAL-3/4 - variable bit rate, delay-tolerant data traffic requiring some sequencing and/or error detection support (originally two AAL types, connection-oriented and connectionless, which have been combined)

AAL-5 - variable bit rate, delay-tolerant, connection-oriented data traffic requiring minimal sequencing or error detection support

ATM Address - Defined in the UNI Specification as 3 formats, each having 20 bytes in length.

ATM Forum - an international non-profit organization formed with the objective of accelerating the use of ATM products and services through a rapid convergence of interoperability specifications. In addition, the Forum promotes industry cooperation and awareness.

 $\textbf{ATM Inverse Multiplexing (AIMUX) -} A \ device \ that \ allows \ multiple \ T1 \ or \ E1 \ communications \ facilities to be combined into a single broadband facility for the transmission of ATM cells.$

ATM Layer link - a section of an ATM Layer connection between two adjacent active ATM Layer entities (ATM-entities).

ATM Link - a virtual path link (VPL) or a virtual channel link (VCL).

ATM Management Interface (AMI) - the user interface to FORE Systems' *ForeThought* switch control software (SCS). AMI lets users monitor and change various operating configurations of FORE Systems switches and network module hardware and software, IP connectivity, and SNMP network management.

ATM Peer-to-Peer Connection - a virtual channel connection (VCC) or a virtual path connection (VPC) directly established, such as workstation-to-workstation. This setup is not commonly used in networks.

ATM Traffic Descriptor - a generic list of parameters that can be used to capture the intrinsic traffic characteristics of a requested ATM connection.

ATM User-to-User Connection - an association established by the ATM Layer to support communication between two or more ATM service users (i.e., between two or more next higher layer entities or between two or more ATM entities). The communication over an ATM Layer connection may be either bidirectional or unidirectional. The same Virtual Channel Identifier (VCI) is used for both directions of a connection at an interface.

atmarp - a FORE program that shows and manipulates ATM ARP entries maintained by the given device driver. This is also used to establish PVC connections.

ATM-attached Host Functional Group (AHFG) - The group of functions performed by an ATM-attached host that is participating in the MPOA service.

atmconfig - a FORE program used to enable or disable SPANS signalling.

atmstat - a FORE program that shows statistics gathered about a given adapter card by the device driver. These statistics include ATM layer and ATM adaptation layer cell and error counts. This can also be used to query other hosts via SNMP.

Attachment User Interface (AUI) - IEEE 802.3 interface between a media attachment unit (MAU) and a network interface card (NIC). The term AUI can also refer to the rear panel port to which an AUI cable might attach.

Auto-logout - a feature that automatically logs out a user if there has been no user interface activity for a specified length of time.

Automatic Protection Switching (APS) - Equipment installed in communications systems to detect circuit failures and automatically switch to redundant, standby equipment.

Available Bit Rate (ABR) - a type of traffic for which the ATM network attempts to meet that traffic's bandwidth requirements. It does not guarantee a specific amount of bandwidth and the end station must retransmit any information that did not reach the far end.

Backbone - the main connectivity device of a distributed system. All systems that have connectivity to the backbone connect to each other, but systems can set up private arrangements with each other to bypass the backbone to improve cost, performance, or security.

Backplane - High-speed communications line to which individual components are connected.

Backward Explicit Congestion Notification (BECN) - A Resource Management cell type generated by the network or the destination, indicating congestion or approaching congestion for traffic flowing in the direction opposite that of the BECN cell.

Bandwidth - usually identifies the capacity or amount of data that can be sent through a given circuit; may be user-specified in a PVC.

Baud - unit of signalling speed, equal to the number of discrete conditions or signal events per second. If each signal event represents only one bit, the baud rate is the same as bps; if each signal event represents more than one bit (such as a dibit), the baud rate is smaller than bps.

Bayonet-Neill-Concelman (BNC) - a bayonet-locking connector used to terminate coaxial cables. BNC is also referred to as Bayonet Network Connector.

Bipolar 8 Zero Substitution (B8ZS) - a technique used to satisfy the ones density requirements of digital T-carrier facilities in the public network while allowing 64 Kbps clear channel data. Strings of eight consecutive zeroes are replaced by an eight-bit code representing two intentional bipolar pulse code violations (000V10V1).

Bipolar Violation (BPV) - an error event on a line in which the normal pattern of alternating high (one) and low (zero) signals is disrupted. A bipolar violation is noted when two high signals occur without an intervening low signal, or vice versa.

B-ISDN Inter-Carrier Interface (B-ICI) - An ATM Forum defined specification for the interface between public ATM networks to support user services across multiple public carriers.

Bit Error Rate (BER) - A measure of transmission quality, generally shown as a negative exponent, (e.g., 10^{-7} which means 1 out of 10^{7} bits [1 out of 10,000,000 bits] are in error).

Bit Interleaved Parity (BIP) - an error-detection technique in which character bit patterns are forced into parity, so that the total number of one bits is always odd or always even. This is accomplished by the addition of a one or zero bit to each byte, as the byte is transmitted; at the other end of the transmission, the receiving device verifies the parity (odd or even) and the accuracy of the transmission.

Bit Robbing - The use of the least significant bit per channel in every sixth frame for signaling.

Bit Stuffing - A process in bit-oriented protocols where a zero is inserted into a string of ones by the sender to prevent the receiver from interpreting valid user data (the string of ones) as control characters (a Flag character for instance).

Border Gateway Protocol (BGP) - used by gateways in an internet connecting autonomous networks. It is derived from experiences learned using the EGP.

bps - bits per second

Bridge - a device that expands a Local Area Network by forwarding frames between data link layers associated with two separate cables, usually carrying a common protocol. Bridges can usually be made to filter certain packets (to forward only certain traffic).

Bridge Protocol Data Unit (BPDU) - A message type used by bridges to exchange management and control information.

Broadband - a service or system requiring transmission channels capable of supporting rates greater than the Integrated Services Digital Network (ISDN) primary rate.

Broadband Access - an ISDN access capable of supporting one or more broadband services.

Broadband Connection Oriented Bearer (BCOB) - Information in the SETUP message that indicates the type of service requested by the calling user.

BCOB-A (Bearer Class A) - Indicated by ATM end user in SETUP message for connection-oriented, constant bit rate service. The network may perform internetworking based on AAL information element (IE).

BCOB-C (Bearer Class C) - Indicated by ATM end user in SETUP message for connection-oriented, variable bit rate service. The network may perform internetworking based on AAL information element (IE).

BCOB-X (Bearer Class X) - Indicated by ATM end user in SETUP message for ATM transport service where AAL, traffic type and timing requirements are transparent to the network.

Broadband Integrated Services Digital Network (B-ISDN) - a common digital network suitable for voice, video, and high-speed data services running at rates beginning at 155 Mbps.

Broadband ISDN User's Part (B-ISUP) - A protocol used to establish, maintain and release broadband switched network connections across an SS7/ATM network.

Broadband Terminal Equipment (B-TE) - An equipment category for B-ISDN which includes terminal adapters and terminals.

Broadcast - Data transmission to all addresses or functions.

Broadcast and Unknown Server (BUS) - in an emulated LAN, the BUS is responsible for accepting broadcast, multicast, and unknown unicast packets from the LECs to the broadcast MAC address (FFFFFFFFFFF) via dedicated point-to-point connections, and forwarding the packets to all of the members of the ELAN using a single point-to-multipoint connection.

Brouter (bridging/router) - a device that routes some protocols and bridges others based on configuration information.

Buffer - A data storage medium used to compensate of a difference in rate of data flow or time of occurrence of events when transmitting data from one device to another.

Building Integrated Timing Supply (BITS) - a master timing supply for an entire building, which is a master clock and its ancillary equipment. The BITS supplies DS1 and/or composite clock timing references for synchronization to all other clocks and timing sources in that building.

Bursty Errored Seconds (BES) - a BES contains more than 1 and fewer than 320 path coding violation error events, and no severely errored frame or AIS defects. Controlled slips are not included in determining BESs.

Bursty Second - a second during which there were at least the set number of BES threshold event errors but fewer than the set number of SES threshold event errors.

Byte - A computer-readable group of bits (normally 8 bits in length).

Call - an association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.

Carrier - a company, such as any of the "baby Bell" companies, that provide network communications services, either within a local area or between local areas.

Carrier Group Alarm (CGA) - A service alarm generated by a channel bank when an out-of-frame (OOF) condition exists for some predetermined length of time (generally 300 milliseconds to 2.5 seconds). The alarm causes the calls using a trunk to be dropped and trunk conditioning to be applied.

Carrier Identification Parameter (CIP) - A 3 or 4 digit code in the initial address message identifying the carrier to be used for the connection.

cchan - a FORE program that manages virtual channels on a FORE switch running asxd.

Cell - an ATM Layer protocol data unit (PDU). The basic unit of information transported in ATM technology, each 53-byte cell contains a 5-byte header and a 48-byte payload.

Cell Delay Variation (CDV) - a quantification of cell clumping for a connection. The cell clumping CDV (yk) is defined as the difference between a cell's expected reference arrival time (ck) and its actual arrival time (ak). The expected reference arrival time (ck) of cell k of a specific connection is max. T is the reciprocal of the negotiated peak cell rate.

Cell Delineation - the protocol for recognizing the beginning and end of ATM cells within the raw serial bit stream.

Cell Header - ATM Layer protocol control information.

Cell Loss Priority (CLP) - the last bit of byte four in an ATM cell header; indicates the eligibility of the cell for discard by the network under congested conditions. If the bit is set to 1, the cell may be discarded by the network depending on traffic conditions.

Cell Loss Ratio - In a network, cell loss ratio is (1-x/y), where y is the number of cells that arrive in an interval at an ingress of the network; and x is the number of these y cells that leave at the egress of the network element.

Cell Loss Ratio (CLR) - CLR is a negotiated QoS parameter and acceptable values are network specific. The objective is to minimize CLR provided the end-system adapts the traffic to the changing ATM layer transfer characteristics. The Cell Loss Ratio is defined for a connection as: Lost Cells/Total Transmitted Cells. The CLR parameter is the value of CLR that the network agrees to offer as an objective over the lifetime of the connection. It is expressed as an order of magnitude, having a range of 10-1 to 10-15 and unspecified.

Cell Misinsertion Rate (CMR) - the ratio of cells received at an endpoint that were not originally transmitted by the source end in relation to the total number of cells properly transmitted.

Cell Rate Adaptation (CRA) - a function performed by a protocol module in which empty cells (known as unassigned cells) are added to the output stream. This is because there always must be a fixed number of cells in the output direction; when there are not enough cells to transmit, unassigned cells are added to the output data stream.

Cell Relay Service (CRS) - a carrier service which supports the receipt and transmission of ATM cells between end users in compliance with ATM standards and implementation specifications.

Cell Transfer Delay - the transit delay of an ATM cell successfully passed between two designated boundaries. See CTD.

Cell Transfer Delay (CTD) - This is defined as the elapsed time between a cell exit event at the measurement point 1 (e.g., at the source UNI) and the corresponding cell entry event at the measurement point 2 (e.g., the destination UNI) for a particular connection. The cell transfer delay between two measurement points is the sum of the total inter-ATM node transmission delay and the total ATM node processing delay.

Channel - A path or circuit along which information flows.

Channel Associated Signaling (CAS) - a form of circuit state signaling in which the circuit state is indicated by one or more bits of signaling status sent repetitively and associated with that specific circuit.

Channel Bank - A device that multiplexes many slow speed voice or data conversations onto high speed link and controls the flow.

Channel Service Unit (CSU) - An interface for digital leased lines which performs loopback testing and line conditioning.

Channelization - capability of transmitting independent signals together over a cable while still maintaining their separate identity for later separation.

Circuit - A communications link between points.

Circuit Emulation Service (CES) - The ATM Forum circuit emulation service interoperability specification specifies interoperability agreements for supporting Constant Bit Rate (CBR) traffic over ATM networks that comply with the other ATM Forum interoperability agreements. Specifically, this specification supports emulation of existing TDM circuits over ATM networks.

Classical IP (CLIP) - IP over ATM which conforms to RFC 1577.

Clear to Send (CTS) - and RS-232 modem interface control signal (sent from the modem to the DTE on pin 5) which indicates that the attached DTE may begin transmitting; issuance in response to the DTE's RTS.

Clocking - Regularly timed impulses.

Closed User Group - A subgroup of network users that can be its own entity; any member of the subgroup can only communicate with other members of that subgroup.

Coaxial Cable - Coax is a type of electrical communications medium used in the LAN environment. This cable consists of an outer conductor concentric to an inner conductor, separated from each other by insulating material, and covered by some protective outer material. This medium offers large bandwidth, supporting high data rates with high immunity to electrical interference and a low incidence of errors. Coax is subject to distance limitations and is relatively expensive and difficult to install.

Cold Start Trap - an SNMP trap which is sent after a power-cycle (see *trap*).

Collision - Overlapping transmissions that occur when two or more nodes on a LAN attempt to transmit at or about the same time.

Committed Information Rate (CIR) - CIR is the information transfer rate which a network offering Frame Relay Services (FRS) is committed to transfer under normal conditions. The rate is averaged over a minimum increment of time.

Common Channel Signaling (CCS) - A form signaling in which a group of circuits share a signaling channel. Refer to SS7.

Common Management Interface Protocol (CMIP) - An ITU-TSS standard for the message formats and procedures used to exchange management information in order to operate, administer maintain and provision a network.

Concatenation - The connection of transmission channels similar to a chain.

Concentrator - a communications device that offers the ability to concentrate many lower-speed channels into and out of one or more high-speed channels.

Configuration - The phase in which the LE Client discovers the LE Service.

Congestion Management - traffic management feature that helps ensure reasonable service for VBR connections in an ATM network, based on a priority, sustained cell rate (SCR), and peak cell rate (PCR). During times of congestion, bandwidth is reduced to the SCR, based on the priority of the connection.

Connection - the concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.

Connection Admission Control (CAC) - the procedure used to decide if a request for an ATM connection can be accepted based on the attributes of both the requested connection and the existing connections.

Connection Endpoint (CE) - a terminator at one end of a layer connection within a SAP.

Connection Endpoint Identifier (CEI) - an identifier of a CE that can be used to identify the connection at a SAP.

Connectionless Broadband Data Service (CBDS) - A connectionless service similar to Bellcore's SMDS defined by European Telecommunications Standards Institute (ETSI).

Connectionless Service - a type of service in which no pre-determined path or link has been established for transfer of information, supported by AAL 4.

Connectionless Service (CLS) - A service which allows the transfer of information among service subscribers without the need for end-to- end establishment procedures.

Connection-Oriented Service - a type of service in which information always traverses the same pre-established path or link between two points, supported by AAL 3.

Constant Bit Rate (CBR) - a type of traffic that requires a continuous, specific amount of bandwidth over the ATM network (e.g., digital information such as video and digitized voice).

Controlled Slip (CS) - a situation in which one frame's worth of data is either lost or replicated. A controlled slip typically occurs when the sending device and receiving device are not using the same clock.

Convergence Sublayer (CS) - a portion of the AAL. Data is passed first to the CS where it is divided into rational, fixed-length packets or PDUs (Protocol Data Units). For example, AAL 4 processes user data into blocks that are a maximum of 64 kbytes long.

Corresponding Entities - peer entities with a lower layer connection among them.

cpath - a FORE program used to manage virtual paths on a FORE switch running asxd.

cport - a FORE program that monitors and changes the state of ports on a FORE switch running asxd.

Cross Connection - a mapping between two channels or paths at a network device.

Customer Premise Equipment (CPE) - equipment that is on the customer side of the point of demarcation, as opposed to equipment that is on a carrier side. See also point of demarcation.

Cut Through - Establishment of a complete path for signaling and/or audio communications.

Cyclic Redundancy Check (CRC) - an error detection scheme in which a number is derived from the data that will be transmitted. By recalculating the CRC at the remote end and comparing it to the value originally transmitted, the receiving node can detect errors.

D3/D4 - Refers to compliance with AT&T TR (Technical Reference) 62411 definitions for coding, supervision, and alarm support. D3/D4 compatibility ensures support of digital PBXes, M24 services, Megacom services, and Mode 3 D3/D4 channel banks at DS-1 level.

D4 Channelization - refers to compliance with AT&T Technical Reference 62411 regarding DS1 frame layout (the sequential assignment of channels and time slot numbers within the DS1).

D4 Framed/Framing Format - in T1, a 193-bit frame format in which the 193rd bit is used for framing and signaling information (the frame/framing bit). To be considered in support of D4 Framing, a device must be able to synchronize and frame-up on the 193rd bit.

Data Communications Equipment (DCE) - a definition in the RS232C standard that describes the functions of the signals and the physical characteristics of an interface for a communication device such as a modem.

Data Country Code (DCC) - This specifies the country in which an address is registered. The codes are given in ISO 3166. The length of this field is two octets. The digits of the data country code are encoded in Binary Coded Decimal (BCD) syntax. The codes will be left justified and padded on the right with the hexadecimal value "F" to fill the two octets.

Data Link - Communications connection used to transmit data from a source to a destination.

Data Link Connection Identifier (DLCI) - connection identifier associated with frame relay packets that serves the same functions as, and translates directly to, the VPI/VCI on an ATM cell.

Data Link Layer - Layer 2 of the OSI model, responsible for encoding data and passing it to the physical medium. The IEEE divides this layer into the LLC (Logical Link Control) and MAC (Media Access Control) sublayers.

Data Set Ready (DSR) - an RS-232 modem interface control signal (sent from the modem to the DTE on pin 6) which indicates that the modem is connected to the telephone circuit. Usually a prerequisite to the DTE issuing RTS.

Data Terminal Equipment (DTE) - generally user devices, such as terminals and computers, that connect to data circuit-terminating equipment. They either generate or receive the data carried by the network.

Data Terminal Ready (DTR) - an RS232 modem interface control signal (sent from the DTE to the modem on pin 20) which indicates that the DTE is ready for data transmission and which requests that the modem be connected to the telephone circuit.

Datagram - a packet of information used in a connectionless network service that is routed to its destination using an address included in the datagram's header.

DECnet - Digital Equipment Corporation's proprietary LAN.

Defense Advanced Research Projects Agency (DARPA) - the US government agency that funded the ARPANET.

Demultiplexing - a function performed by a layer entity that identifies and separates SDUs from a single connection to more than one connection (see *multiplexing*).

Destination End Station (DES) - An ATM termination point which is the destination for ATM messages of a connection and is used as a reference point for ABR services. See SES.

Digital Access and Cross-Connect System (DACS) - Digital switching system for routing T1 lines, and DS-0 portions of lines, among multiple T1 ports.

Digital Cross-connect System (DCS) - an electronic patch panel used to route digital signals in a central office.

Digital Standard n (0, 1, 1C, 2, and 3) (DSn) - a method defining the rate and format of digital hierarchy, with asynchronous data rates defined as follows:

DS0	64kb/s	1 voice channel
DS1	1.544Mb/s	24 DS0s
DS1C	3.152 Mb/s	2 DS1s
DS2	6.312 Mb/s	4 DS1s
DS3	44.736 Mb/s	28 DS1s

Synchronous data rates (SONET) are defined as:

STS-1/OC-1	51.84 Mb/s	28 DS1s or 1 DS3
STS-3/OC-3	155.52 Mb/s	3 STS-1s byte interleaved
STS-3c/OC-3c	155.52 Mb/s	Concatenated, indivisible payload
STS-12/OC-12	622.08 Mb/s	12 STS-1s, 4 STS-3cs, or any mixture
STS-12c/OC-12c	622.08 Mb/s	Concatenated, indivisible payload
STS-48/OC-48	2488.32 Mb/s	48 STS-1s, 16 STS-3cs, or any mixture

DIP (Dual In-line Package) Switch - a device that has two parallel rows of contacts that let the user switch electrical current through a pair of those contacts to on or off. They are used to reconfigure components and peripherals.

Domain Name Server - a computer that converts names to their corresponding Internet numbers. It allows users to telnet or FTP to the name instead of the number.

Domain Naming System (DNS) - the distributed name and address mechanism used in the Internet.

Duplex - Two way communication.

DXI - a generic phrase used in the full names of several protocols, all commonly used to allow a pair of DCE and DTE devices to share the implementation of a particular WAN protocol. The protocols define the packet formats used to transport data between DCE and DTE devices.

DXI Frame Address (DFA) - a connection identifier associated with ATM DXI packets that serves the same functions as, and translates directly to, the VPI/VCI on an ATM cell.

Dynamic Allocation - A technique in which the resources assigned for program execution are determined by criteria applied at the moment of need.

E.164 - A public network addressing standard utilizing up to a maximum of 15 digits. ATM uses E.164 addressing for public network addressing.

E1 - Wide-area digital transmission scheme used predominantly in Europe that carries data at a rate of 2.048 Mbps. E1 lines can be leased for private use from common carriers.

E3 - Wide-area digital transmission scheme used predominantly in Europe that carries data at a rate of 34.368 Mbps. E3 lines can be leased for private use from common carriers.

Early Packet Discard - A mechanism that allows a complete AAL5 frame to be discarded when a threshold condition is met (e.g., one indicating imminent congestion). EPD helps to avoid congestion that jeopardizes a switch's ability to support existing connections with a guaranteed service.

Edge Device - A physical device which is capable of forwarding packets between legacy interworking interfaces (e.g., Ethernet, Token Ring, etc.) and ATM interfaces based on data-link and network layer information but which does not participate in the running of any network layer routing protocol. An Edge Device obtains forwarding descriptions using the route distribution protocol.

elarp - a FORE program that shows and manipulates MAC and ATM address mappings for LAN Emulation Clients (LECs).

elconfig - a FORE program that shows and modifies LEC configuration. Lets the user set the NSAP address of the LAN Emulation Configuration Server, display the list of Emulated LANs configured in the LECS for this host, display the list of ELANs locally configured along with the membership state of each, and locally administer ELAN membership.

Electrically Erasable Programmable Read Only Memory (EEPROM) - an EPROM that can be cleared with electrical signals rather than the traditional ultraviolet light.

Electromagnetic Interference (EMI) - signals generated and radiated by an electronic device that cause interference with radio communications, among other effects.

Electronics Industries Association (EIA) - a USA trade organization that issues its own standards and contributes to ANSI; developed RS-232. Membership includes USA manufacturers.

Embedded SNMP Agent - an SNMP agent can come in two forms: embedded or proxy. An embedded SNMP agent is integrated into the physical hardware and software of the unit.

Emulated Local Area Network (ELAN) - A logical network initiated by using the mechanisms defined by LAN Emulation. This could include ATM and legacy attached end stations.

End System (ES) - a system where an ATM connection is terminated or initiated (an originating end system initiates the connection).

End System Identifier (ESI) - This identifier distinguishes multiple nodes at the same level in case the lower level peer group is partitioned.

End-to-End Connection - when used in reference to an ATM network, a connection that travels through an ATM network, passing through various ATM devices and with endpoints at the termination of the ATM network.

Enterprise - Terminology generally referring to customers with multiple, non-contiguous geographic locations.

Equalization (EQL) - the process of compensating for line distortions.

Erasable Programmable Read Only Memory (EPROM) - A PROM which may be erased and rewritten to perform new or different functions (normally done with a PROM burner).

Errored Second (ES) - a second during which at least one code violation occurred.

Ethernet - a 10-Mbps, coaxial standard for LANs in which all nodes connect to the cable where they contend for access.

Excessive Zeroes (EXZ) Error Event - An Excessive Zeroes error event for an AMI-coded signal is the occurrence of more than fifteen contiguous zeroes. For a B8ZS coded signal, the defect occurs when more than seven contiguous zeroes are detected.

Explicit Forward Congestion Indication (EFCI) - the second bit of the payload type field in the header of an ATM cell, the EFCI bit indicates network congestion to receiving hosts. On a congested switch, the EFCI bit is set to "1" by the transmitting network module when a certain number of cells have accumulated in the network module's shared memory buffer. When a cell is received that has its EFCI bit set to "1," the receiving host notifies the sending host, which should then reduce its transmission rate.

Explicit Rate (ER) - The Explicit Rate is an RM-cell field used to limit the source ACR to a specific value. It is initially set by the source to a requested rate (such as PCR). It may be subsequently reduced by any network element in the path to a value that the element can sustain. ER is formatted as a rate.

Extended Industry Standard Architecture (EISA) - bus architecture for desktop computers that provides a 32-bit data passage and maintains compatibility with the ISA or AT architecture.

Extended Super Frame (ESF) - a T1 framing format that utilizes the 193rd bit as a framing bit, but whose Superframe is made up of 24 frames instead of 12 as in D4 format. ESF also provides CRC error detection and maintenance data link functions.

Exterior Gateway Protocol (EGP) - used by gateways in an internet, connecting autonomous networks.

Fairness - related to Generic Flow Control, fairness is defined as meeting all of the agreed quality of service requirements by controlling the order of service for all active connections.

Far End Block Error (FEBE) - an error detected by extracting the 4-bit FEBE field from the path status byte (G1). The legal range for the 4-bit field is between 0000 and 1000, representing zero to eight errors. Any other value is interpreted as zero errors.

Far End Receive Failure (FERF) - a line error asserted when a 110 binary pattern is detected in bits 6, 7, 8 of the K2 byte for five consecutive frames. A line FERF is removed when any pattern other than 110 is detected in these bits for five consecutive frames.

Far-End - in a relationship between two devices in a circuit, the far-end device is the one that is remote.

Face Contact (FC) - Designation for fiber optic connector designed by Nippon Telegraph and Telephone which features a movable anti-rotation key allowing good repeatable performance despite numerous mating. Normally referred to as Fiber Connector, FC actually stands for Face Contact and sometimes linked with PC (Point Contact), designated as FC or FC-PC.

FCC Part 68 - The FCC rules regulating the direct connection of non-telephone company provided equipment to the public telephone network.

Federal Communications Commission (FCC) - a board of commissioners appointed by the President under the Communications Act of 1934, with the authority to regulate all interstate telecommunications originating in the United States, including transmission over phone lines.

Fiber Distributed Data Interface (FDDI) - high-speed data network that uses fiber-optic as the physical medium. Operates in similar manner to Ethernet or Token Ring, only faster.

File Transfer Protocol (FTP) - a TCP/IP protocol that lets a user on one computer access, and transfer data to and from, another computer over a network. ftp is usually the name of the program the user invokes to accomplish this task.

First-In, First-Out (FIFO) - method of coordinating the sequential flow of data through a buffer.

Flag - a bit pattern of six binary "1"s bounded by a binary "0" at each end (forms a 0111 1110 or Hex "7E"). It is used to mark the beginning and/or end of a frame.

Flow Control - The way in which information is controlled in a network to prevent loss of data when the receiving buffer is near its capacity.

ForeThought PNNI (FT-PNNI) - a FORE Systems routing and signalling protocol that uses private ATM (NSAP) addresses; a precursor to ATM Forum PNNI (see PNNI).

Forward Error Correction (FEC) - A technique used by a receiver for correcting errors incurred in transmission over a communications channel without requiring retransmission of any information by the transmitter; typically involves a convolution of the transmitted bits and the appending of extra bits by both the receiver and transmitter using a common algorithm.

Forward Explicit Congestion Notification (FECN) - Bit set by a Frame Relay network to inform data terminal equipment (DTE) receiving the frame that congestion was experienced in the path from source to destination. DTE receiving frames with the FECN bit set can request that higher-level protocols take flow control action as appropriate.

Fractional T1 - the use of bandwidth in 64Kbps increments up to 1.544Mbps from a T1 facility.

Frame - a variable length group of data bits with a specific format containing flags at the beginning and end to provide demarcation.

Frame Check Sequence (FCS) - In bit-oriented protocols, a 16-bit field that contains transmission error checking information, usually appended to the end of the frame.

Frame Relay - a fast packet switching protocol based on the LAPD protocol of ISDN that performs routing and transfer with less overhead processing than X.25.

Frame Synchronization Error - an error in which one or more time slot framing bits are in error.

Frame-Based UNI (FUNI) - An ATM switch-based interface which accepts frame-based ATM traffic and converts it into cells.

Frame-Relay Service (FRS) - A connection oriented service that is capable of carrying up to 4096 bytes per frame.

Framing - a protocol that separates incoming bits into identifiable groups so that the receiving multiplexer recognizes the grouping.

Frequency Division Multiplexing (FDM) - a method of dividing an available frequency range into parts with each having enough bandwidth to carry one channel.

Gbps - gigabits per second (billion)

Generic Cell Rate Algorithm (GCRA) - an algorithm which is employed in traffic policing and is part of the user/network service contract. The GCRA is a scheduling algorithm which ensures that cells are marked as conforming when they arrive when expected or later than expected and non-conforming when they arrive sooner than expected.

Generic Connection Admission Control (GCAC) - This is a process to determine if a link has potentially enough resources to support a connection.

Generic Flow Control (GFC) - the first four bits of the first byte in an ATM cell header. Used to control the flow of traffic across the User-to-Network Interface (UNI), and thus into the network. Exact mechanisms for flow control are still under investigation and no explicit definition for this field exists at this time. (This field is used only at the UNI; for NNI-NNI use (between network nodes), these four bits provide additional network address capacity, and are appended to the VPI field.)

GIO - a proprietary bus architecture used in certain Silicon Graphics, Inc. workstations.

Header - protocol control information located at the beginning of a protocol data unit.

Header Error Control (HEC) - a CRC code located in the last byte of an ATM cell header that is used for checking cell header integrity only.

High Density Bipolar (HDB3) - A bipolar coding method that does not allow more than 3 consecutive zeroes.

High Level Data Link Control (HDLC) - An ITU-TSS link layer protocol standard for point-to-point and multi-point communications.

High Performance Parallel Interface (HIPPI) - ANSI standard that extends the computer bus over fairly short distances at speeds of 800 and 1600 Mbps.

High-Speed Serial Interface (HSSI) - a serial communications connection that operates at speeds of up to 1.544 Mbps.

Host - In a network, the primary or controlling computer in a multiple computer installation.

HPUX - the Hewlett-Packard version of UNIX.

Hub - a device that connects several other devices, usually in a star topology.

I/O Module - FORE's interface cards for the LAX-20 LAN Access Switch, designed to connect Ethernet, Token Ring, and FDDI LANs to FORE ATM networks.

Institute of Electrical and Electronics Engineers (IEEE) - the world's largest technical professional society. Based in the U.S., the IEEE sponsors technical conferences, symposia & local meetings worldwide, publishes nearly 25% of the world's technical papers in electrical, electronics & computer engineering, provides educational programs for members, and promotes standardization.

IEEE 802 - Standards for the interconnection of LAN computer equipment. Deals with the Data Link Layers of the ISO Reference Model for OSI.

IEEE 802.1 - Defines the high-level network interfaces such as architecture, internetworking and network management.

IEEE 802.2 - Defines the Logical Link Control interface between the Data Link and Network Layers.

IEEE 802.3 - Defines CSMA/CD (Ethernet).

IEEE 802.4 - Defines the token-passing bus.

IEEE 802.5 - Defines the Token Ring access methodology. This standard incorporates IBM's Token Ring specifications.

IEEE 802.6 - Defines Metropolitan Area Networks.

IEEE 802.8 - The fiber optics technical advisory group.

IEEE 802.9 - Defines integrated data and voice networks.

Integrated Services Digital Network (ISDN) - an emerging technology that is beginning to be offered by the telephone carriers of the world. ISDN combines voice and digital network services into a single medium or wire.

Interexchange Carriers (IXC) - Long-distance communications companies that provide service between Local Access Transport Areas (LATAs).

Interface Data - the unit of information transferred to/from the upper layer in a single interaction across a SAP. Each Interface Data Unit (IDU) controls interface information and may also contain the whole or part of the SDU.

Interface Data Unit (IDU) - The unit of information transferred to/from the upper layer in a single interaction across the SAP. Each IDU contains interface control information and may also contain the whole or part of the SDU.

Interim Local Management Interface (ILMI) - the standard that specifies the use of the Simple Network Management Protocol (SNMP) and an ATM management information base (MIB) to provide network status and configuration information.

Intermediate System (IS) - a system that provides forwarding functions or relaying functions or both for a specific ATM connection. OAM cells may be generated and received.

International Standards Organization (ISO) - a voluntary, non treaty organization founded in 1946 that is responsible for creating international standards in many areas, including computers and communications.

International Telephone and Telegraph Consultative Committee (CCITT) - the international standards body for telecommunications.

Internet - (note the capital "I") the largest internet in the world including large national backbone nets and many regional and local networks worldwide. The Internet uses the TCP/IP suite. Networks with only e-mail connectivity are not considered on the Internet.

internet - while an internet is a network, the term "internet" is usually used to refer to a collection of networks interconnected with routers.

Internet Addresses - the numbers used to identify hosts on an internet network. Internet host numbers are divided into two parts; the first is the network number and the second, or local, part is a host number on that particular network. There are also three classes of networks in the Internet, based on the number of hosts on a given network. Large networks are classified as Class A, having addresses in the range 1-126 and having a maximum of 16,387,064 hosts. Medium networks are classified as Class B, with addresses in the range 128-191 and with a maximum of 64,516 hosts. Small networks are classified as Class C, having addresses in the range 192-254 with a maximum of 254 hosts. Addresses are given as dotted decimal numbers in the following format:

nnn.nnn.nnn.nnn

In a Class A network, the first of the numbers is the network number, the last three numbers are the local host address.

In a Class B network, the first two numbers are the network, the last two are the local host address.

In a Class C network, the first three numbers are the network address, the last number is the local host address.

The following table summarizes the classes and sizes:

Glossary

Class	First #	Max# Hosts
A	1-126	16,387,064
В	129-191	64,516
C	192-223	254

Network mask values are used to identify the network portion and the host portion of the address. Default network masks are as follows:

Class A - 255.0.0.0

Class B - 255.255.0.0

Class C - 255.255.255.0

Subnet masking is used when a portion of the host ID is used to identify a subnetwork. For example, if a portion of a Class B network address is used for a subnetwork, the mask could be set as 255.255.255.0. This would allow the third byte to be used as a subnetwork address. All hosts on the network would still use the IP address to get on the Internet.

Internet Control Message Protocol (ICMP) - the protocol that handles errors and control messages at the IP layer. ICMP is actually a part of the IP protocol layer. It can generate error messages, test packets, and informational messages related to IP.

Internet Engineering Task Force (IETF) - a large, open, international community of network designers, operators, vendors and researchers whose purpose is to coordinate the operation, management and evolution of the Internet to resolve short- and mid-range protocol and architectural issues.

Internet Protocol (IP) - a connectionless, best-effort packet switching protocol that offers a common layer over dissimilar networks.

Internetwork Packet Exchange (IPX) Protocol - a NetWare protocol similar to the Xerox Network Systems (XNS) protocol that provides datagram delivery of messages.

Interoperability - The ability of software and hardware on multiple machines, from multiple vendors, to communicate.

Interworking Function (IWF) - provides a means for two different technologies to interoperate.

IP Address - a unique 32-bit integer used to identify a device in an IP network. You will most commonly see IP addresses written in "dot" notation (e.g., 192.228.32.14).

IP Netmask - a 32-bit pattern that is combined with an IP address to determine which bits of an IP address denote the network number and which denote the host number. Netmasks are useful for sub-dividing IP networks. IP netmasks are written in "dot" notation (e.g., 255.255.0.0).

ISA Bus - a bus standard developed by IBM for expansion cards in the first IBM PC. The original bus supported a data path only 8 bits wide. IBM subsequently developed a 16-bit version for its AT class computers. The 16-bit AT ISA bus supports both 8- and 16-bit cards. The 8-bit bus is commonly called the PC/XT bus, and the 16-bit bus is called the AT bus.

Isochronous - signals carrying embedded timing information or signals that are dependent on uniform timing; usually associated with voice and/or video transmission.

International Telecommunications Union Telecommunications (ITU-T) - an international body of member countries whose task is to define recommendations and standards relating to the international telecommunications industry. The fundamental standards for ATM have been defined and published by the ITU-T (Previously CCITT).

J2 - Wide-area digital transmission scheme used predominantly in Japan that carries data at a rate of 6.312 Mbps.

Jitter - analog communication line distortion caused by variations of a signal from its reference timing position.

Joint Photographic Experts Group (JPEG) - An ISO Standards group that defines how to compress still pictures.

Jumper - a patch cable or wire used to establish a circuit, often temporarily, for testing or diagnostics; also, the devices, shorting blocks, used to connect adjacent exposed pins on a printed circuit board that control the functionality of the card.

Kbps - kilobits per second (thousand)

LAN Access Concentrator - a LAN access device that allows a shared transmission medium to accommodate more data sources than there are channels currently available within the transmission medium.

LAN Emulation Address Resolution Protocol (LE_ARP) - A message issued by a LE client to solicit the ATM address of another function.

LAN Emulation Client (LEC) - the component in an end system that performs data forwarding, address resolution, and other control functions when communicating with other components within an ELAN.

LAN Emulation Configuration Server (LECS) - the LECS is responsible for the initial configuration of LECs. It provides information about available ELANs that a LEC may join, together with the addresses of the LES and BUS associated with each ELAN.

LAN Emulation Server (LES) - the LES implements the control coordination function for an ELAN by registering and resolving MAC addresses to ATM addresses.

LAN Emulation (LANE) - technology that allows an ATM network to function as a LAN backbone. The ATM network must provide multicast and broadcast support, address mapping (MAC-to-ATM), SVC management, and a usable packet format. LANE also defines Ethernet and Token Ring ELANs.

lane - a program that provides control over the execution of the LAN Emulation Server (LES), Broadcast/Unknown Server (BUS), and LAN Emulation Configuration Server (LECS) on the local host.

Latency - The time interval between a network station seeking access to a transmission channel and that access being granted or received.

Layer Entity - an active layer within an element.

Layer Function - a part of the activity of the layer entities.

Layer Service - a capability of a layer and the layers beneath it that is provided to the upper layer entities at the boundary between that layer and the next higher layer.

Layer User Data - the information transferred between corresponding entities on behalf of the upper layer or layer management entities for which they are providing services.

le - a FORE program that implements both the LAN Emulation Server (LES) and the Broadcast/Unknown Server (BUS).

Leaky Bucket - informal cell policing term for the Generic Cell Rate Algorithm which in effect receives cells into a bucket and leaks them out at the specified or contracted rate (i.e., PCR).

Least Significant Bit (LSB) - lowest order bit in the binary representation of a numerical value.

lecs - a FORE program that implements the assignment of individual LECs to different emulated LANs.

leq - a FORE program that provides information about an ELAN. This information is obtained from the LES, and includes MAC addresses registered on the ELAN together with their corresponding ATM addresses.

Line Build Out (LBO) - Because T1 circuits require the last span to lose 15-22.5 dB, a selectable output attenuation is generally required of DTE equipment (typical selections include 0.0, 7.5 and 15 dB of loss at 772 KHz).

Line Code Violations (LCV) - Error Event. A Line Coding Violation (LCV) is the occurrence of either a Bipolar Violation (BPV) or Excessive Zeroes (EXZ) Error Event.

Link - An entity that defines a topological relationship (including available transport capacity) between two nodes in different subnetworks. Multiple links may exist between a pair of subnetworks. Synonymous with logical link.

Link Access Procedure, Balanced (LAPB) - Data link protocol in the X.25 protocol stack. LAPB is a bit-oriented protocol derived from HDLC. See also HDLC and X.25.

Link Down Trap - an SNMP trap, sent when an interface changes from a normal state to an error state, or is disconnected.

Link Layer - layer in the OSI model regarding transmission of data between network nodes.

Link Up Trap - an SNMP trap, sent when an interface changes from an error condition to a normal state.

Load Sharing - Two or more computers in a system that share the load during peak hours. During periods of non peak hours, one computer can manage the entire load with the other acting as a backup.

Local Access and Transport Area (LATA) - Geographic boundaries of the local telephone network, specified by the FCC, in which a single LEC may perform its operations. Communications outside or between LATAs are provided by IXCs.

Local Area Network (LAN) - a data network intended to serve an area of only a few square kilometers or less. Because the network is known to cover only a small area, optimizations can be made in the network signal protocols that permit higher data rates.

Logical Link Control (LLC) - protocol developed by the IEEE 802 committee for data-link-layer transmission control; the upper sublayer of the IEEE Layer 2 (OSI) protocol that complements the MAC protocol; IEEE standard 802.2; includes end-system addressing and error checking.

Loopback - a troubleshooting technique that returns a transmitted signal to its source so that the signal can be analyzed for errors. Typically, a loopback is set at various points in a line until the section of the line that is causing the problem is discovered.

looptest - program that tests an interface for basic cell reception and transmission functionality, usually used for diagnostic purposes to determine if an interface is functioning properly.

Loss Of Frame (LOF) - a type of transmission error that may occur in wide-area carrier lines.

Loss Of Pointer (LOP) - a type of transmission error that may occur in wide-area carrier lines.

Loss Of Signal (LOS) - a type of transmission error that may occur in wide-area carrier lines, or a condition declared when the DTE senses a loss of a DS1 signal from the CPE for more the 150 milliseconds (the DTE generally responds with an all ones "Blue or AIS" signal).

Management Information Base (MIB) - the set of parameters that an SNMP management station can query or set in the SNMP agent of a networked device (e.g., router).

Maximum Burst Size (MBS) - the Burst Tolerance (BT) is conveyed through the MBS which is coded as a number of cells. The BT together with the SCR and the GCRA determine the MBS that may be transmitted at the peak rate and still be in conformance with the GCRA.

Maximum Burst Tolerance - the largest burst of data that a network device is guaranteed to handle without discarding cells or packets. Bursts of data larger than the maximum burst size may be subject to discard.

Maximum Cell Delay Variance (MCDV) - This is the maximum two-point CDV objective across a link or node for the specified service category.

Maximum Cell Loss Ratio (MCLR) - This is the maximum ratio of the number of cells that do not make it across the link or node to the total number of cells arriving at the link or node.

Maximum Cell Transfer Delay (MCTD) - This is the sum of the fixed delay component across the link or node and MCDV.

 $\label{lem:maximum Transmission Unit (MTU) - the largest unit of data that can be sent over a type of physical medium.$

Mbps - megabits per second (million)

Media Access Control (MAC) - a media-specific access control protocol within IEEE 802 specifications; currently includes variations for Token Ring, token bus, and CSMA/CD; the lower sublayer of the IEEE's link layer (OSI), which complements the Logical Link Control (LLC).

Media Attachment Unit (MAU) - device used in Ethernet and IEEE 802.3 networks that provides the interface between the AUI port of a station and the common medium of the Ethernet. The MAU, which can be built into a station or can be a separate device, performs physical layer functions including conversion of the digital data from the Ethernet interface, collision detection, and injection of bits onto the network.

Media Interface Connector (MIC) - fiber optic connector that joins fiber to the FDDI controller.

Message Identifier (MID) - message identifier used to associate ATM cells that carry segments from the same higher layer packet.

Metasignalling - an ATM Layer Management (LM) process that manages different types of signalling and possibly semipermanent virtual channels (VCs), including the assignment, removal, and checking of VCs.

Metasignalling VCs - the standardized VCs that convey metasignalling information across a User-to-Network Interface (UNI).

Metropolitan Area Network (MAN) - network designed to carry data over an area larger than a campus such as an entire city and its outlying area.

MicroChannel - a proprietary 16- or 32-bit bus developed by IBM for its PS/2 computers' internal expansion cards; also offered by others.

Minimum Cell Rate (MCR) - parameter defined by the ATM Forum for ATM traffic management, defined only for ABR transmissions and specifying the minimum value for the ACR.

Most Significant Bit (MSB) - highest order bit in the binary representation of a numerical value.

Motion Picture Experts Group (MPEG) - ISO group dealing with video and audio compression techniques and mechanisms for multiplexing and synchronizing various media streams.

MPOA Client - A device which implements the client side of one or more of the MPOA protocols, (i.e., is a SCP client and/or an RDP client. An MPOA Client is either an Edge Device Functional Group (EDFG) or a Host Behavior Functional Group (HBFG).

MPOA Server - An MPOA Server is any one of an ICFG or RSFG.

MPOA Service Area - The collection of server functions and their clients. A collection of physical devices consisting of an MPOA server plus the set of clients served by that server.

MPOA Target - A set of protocol address, path attributes, (e.g., internetwork layer QoS, other information derivable from received packet) describing the intended destination and its path attributes that MPOA devices may use as lookup keys.

Mu-Law - The PCM coding and companding standard used in Japan and North America.

Multicasting - The ability to broadcast messages to one node or a select group of nodes.

Multi-homed - a device having both an ATM and another network connection, like Ethernet.

Multimode Fiber Optic Cable (MMF) - fiber optic cable in which the signal or light propagates in multiple modes or paths. Since these paths may have varying lengths, a transmitted pulse of light may be received at different times and smeared to the point that pulses may interfere with surrounding pulses. This may cause the signal to be difficult or impossible to receive. This pulse dispersion sometimes limits the distance over which a MMF link can operate.

Multiplexing - a function within a layer that interleaves the information from multiple connections into one connection (see demultiplexing).

Multipoint Access - user access in which more than one terminal equipment (TE) is supported by a single network termination.

Multipoint-to-Multipoint Connection - a collection of associated ATM VC or VP links, and their associated endpoint nodes, with the following properties:

- 1. All N nodes in the connection, called Endpoints, serve as a Root Node in a Point-to-Multipoint connection to all of the (N-1) remaining endpoints.
- 2. Each of the endpoints can send information directly to any other endpoint, but the receiving endpoint cannot distinguish which of the endpoints is sending information without additional (e.g., higher layer) information.

Multipoint-to-Point Connection - a Point-to-Multipoint Connection may have zero bandwidth from the Root Node to the Leaf Nodes, and non-zero return bandwidth from the Leaf Nodes to the Root Node. Such a connection is also known as a Multipoint-to-Point Connection.

Multiprotocol over ATM (MPOA) - An effort taking place in the ATM Forum to standardize protocols for the purpose of running multiple network layer protocols over ATM.

Narrowband Channel - sub-voicegrade channel with a speed range of 100 to 200 bps.

National TV Standards Committee (NTSC) - Started in the US in 1953 from a specification laid down by the National Television Standards Committee. It takes the B-Y and R-Y color difference signals, attenuates them to I and Q, then modulates them using double-sideband suppressed subcarrier at 3.58MHz. The carrier reference is sent to the receiver as a burst during the back porch. An industry group that defines how television signals are encoded and transmitted in the US. (See also PAL, SECAM for non-U.S. countries).

Near-End - in a relationship between two devices in a circuit, the near-end device is the one that is local.

Network Layer - Layer three In the OSI model, the layer that is responsible for routing data across the network.

Network Management Entity (NM) - body of software in a switching system that provides the ability to manage the PNNI protocol. NM interacts with the PNNI protocol through the MIB.

Network Management Layer (NML) - an abstraction of the functions provided by systems which manage network elements on a collective basis, providing end-to-end network monitoring.

Network Management Station (NMS) - system responsible for managing a network or portion of a network by talking to network management agents, which reside in the managed nodes.

Network Module - ATM port interface cards which may be individually added to or removed from any FORE ATM switch to provide a diverse choice of connection alternatives.

Network Parameter Control (NPC) - Defined as the set of actions taken by the network to monitor and control traffic from the NNI. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior which can affect the QoS of other already established connections by detecting violations of negotiated parameters and taking appropriate actions. Refer to UPC.

Network Redundancy - Duplicated network equipment and/or data which can provide a backup in case of network failures.

Network Service Access Point (NSAP) - OSI generic standard for a network address consisting of 20 octets. ATM has specified E.164 for public network addressing and the NSAP address structure for private network addresses.

Network-to-Network Interface or Network Node Interface (NNI) - the interface between two public network pieces of equipment.

Node - A computer or other device when considered as part of a network.

Non Return to Zero (NRZ) - a binary encoding scheme in which ones and zeroes are represented by opposite and alternating high and low voltages and where there is no return to a zero (reference) voltage between encoded bits.

Non Return to Zero Inverted (NRZI) - A binary encoding scheme that inverts the signal on a "1" and leaves the signal unchanged for a "0". (Also called transition encoding.)

Nonvolatile Storage - Memory storage that does not lose its contents when power is turned off.

NuBus - a high-speed bus used in Macintosh computers, structured so users can put a card into any slot on the board without creating conflict over the priority between those cards.

nx64K - This refers to a circuit bandwidth or speed provided by the aggregation of nx64 kbps channels (where n= integer > 1). The 64K or DS0 channel is the basic rate provided by the T Carrier systems.

Nyquist Theorem - In communications theory, a formula stating that two samples per cycle is sufficient to characterize a bandwidth limited analog signal; in other words, the sampling rate must be twice the highest frequency component of the signal (i.e., sample 4 KHz analog voice channels 8000 times per second).

Object Identifier (OID) - the address of a MIB variable.

 $\textbf{Octet -} \ a \ grouping \ of \ 8 \ bits; \ similar, \ but \ not \ identical \ to, \ a \ byte.$

One's Density - The requirement for digital transmission lines in the public switched telephone network that eight consecutive "0"s cannot be in a digital data stream; exists because repeaters and clocking devices within the network will lose timing after receiving eight "0"s in a row; a number of techniques are used to insert a "1" after every seventh-consecutive "0" (see Bit Stuffing).

Open Shortest Path First (OSPF) Protocol - a routing algorithm for IP that incorporates least-cost, equal-cost, and load balancing.

Open Systems Interconnection (OSI) - the 7-layer suite of protocols designed by ISO committees to be the international standard computer network architecture.

OpenView - Hewlett-Packard's network management software.

Operation and Maintenance (OAM) Cell - a cell that contains ATM LM information. It does not form part of the upper layer information transfer.

Optical Carrier level-n (OC-n) - The optical counterpart of STS-n (the basic rate of 51.84 Mbps on which SONET is based is referred to as OC-1 or STS-1).

Organizationally Unique Identifier (OUI) - Part of RFC 1483. A three-octet field in the SubNetwork Attachment Point (SNAP) header, identifying an organization which administers the meaning of the following two octet Protocol Identifier (PID) field in the SNAP header. Together they identify a distinct routed or bridged protocol.

Out-of-Band Management - refers to switch configuration via the serial port or over Ethernet, not ATM.

Out-of-Frame (OOF) - a signal condition and alarm in which some or all framing bits are lost.

Packet - An arbitrary collection of data grouped and transmitted with its user identification over a shared facility.

Packet Assembler Disassembler (PAD) - interface device that buffers data sent to/from character mode devices, and assembles and disassembles the packets needed for X.25 operation.

Packet Internet Groper (ping) - a program used to test reachability of destinations by sending them an ICMP echo request and waiting for a reply.

Packet Level Protocol (PLP) - network layer protocol in the X.25 protocol stack. Sometimes called X.25 Level 3 or X.25 Protocol.

Packet Switched Network (PSN) - a network designed to carry data in the form of packets. The packet and its format is internal to that network.

Packet Switching - a communications paradigm in which packets (messages) are individually routed between hosts with no previously established communications path.

Partial Packet Discard - a mechanism that allows the remaining cells of an AAL5 frame to be discarded when one or more cells of that AAL5 frame have been dropped. PPD eliminates the sending of partial AAL5 frames through the ATM network that will have to be retransmitted by the sender.

Payload Scrambling - a technique that eliminates certain bit patterns that may occur within an ATM cell payload that could be misinterpreted by certain sensitive transmission equipment as an alarm condition.

Payload Type (PT) - bits 2...4 in the fourth byte of an ATM cell header. The PT indicates the type of information carried by the cell. At this time, values 0...3 are used to identify various types of user data, values 4 and 5 indicate management information, and values 6 and 7 are reserved for future use.

Peak Cell Rate - at the PHY Layer SAP of a point-to-point VCC, the Peak Cell Rate is the inverse of the minimum inter-arrival time T0 of the request to send an ATM-SDU.

Peak Cell Rate (PCR) - parameter defined by the ATM Forum for ATM traffic management. In CBR transmissions, PCR determines how often data samples are sent. In ABR transmissions, PCR determines the maximum value of the ACR.

Peer Entities - entities within the same layer.

Peripheral Component Interconnect (PCI) - a local-bus standard created by Intel.

Permanent Virtual Channel Connection (PVCC) - A Virtual Channel Connection (VCC) is an ATM connection where switching is performed on the VPI/VCI fields of each cell. A Permanent VCC is one which is provisioned through some network management function and left up indefinitely.

Permanent Virtual Circuit (or Channel) (PVC) - a circuit or channel through an ATM network provisioned by a carrier between two endpoints; used for dedicated long-term information transport between locations.

Permanent Virtual Path Connection (PVPC) - A Virtual Path Connection (VPC) is an ATM connection where switching is performed on the VPI field only of each cell. A PVPC is one which is provisioned through some network management function and left up indefinitely.

Phase Alternate Line (PAL) - Largely a German/British development in the late 60s, used in the UK and much of Europe. The B-Y and R-Y signals are weighted to U and V, then modulated onto a double-sideband suppressed subcarrier at 4.43MHz. The V (R-Y) signal's phase is turned through 180 degrees on each alternate line. This gets rid of NTSC's hue changes with phase errors at the expense of de-saturation. The carrier reference is sent as a burst in the back porch. The phase of the burst is alternated every line to convey the phase switching of the V signal. The burst's average phase is -V. (see NTSC for U.S.).

Physical Layer (PHY) - the actual cards, wires, and/or fiber-optic cabling used to connect computers, routers, and switches.

Physical Layer Connection - an association established by the PHY between two or more ATM-entities. A PHY connection consists of the concatenation of PHY links in order to provide an end-to-end transfer capability to PHY SAPs.

Physical Layer Convergence Protocol (PLCP) - a framing protocol that runs on top of the T1 or E1 framing protocol.

Physical Medium (PM) - Refers to the actual physical interfaces. Several interfaces are defined including STS-1, STS-3c, STS-12c, STM-1, STM-4, DS1, E1, DS2, E3, DS3, E4, FDDI-based, Fiber Channel-based, and STP. These range in speeds from 1.544Mbps through 622.08 Mbps.

Physical Medium Dependent (PMD) - a sublayer concerned with the bit transfer between two network nodes. It deals with wave shapes, timing recovery, line coding, and electro-optic conversions for fiber based links.

Plesiochronous - two signals are plesiochronous if their corresponding significant instants occur at nominally the same rate, with variations in rate constrained to specified limits.

Point of Demarcation - the dividing line between a carrier and the customer premise that is governed by strict standards that define the characteristics of the equipment on each side of the demarcation. Equipment on one side of the point of demarcation is the responsibility of the customer. Equipment on the other side of the point of demarcation is the responsibility of the carrier.

Point-to-Multipoint Connection - a collection of associated ATM VC or VP links, with associated endpoint nodes, with the following properties:

- 1. One ATM link, called the Root Link, serves as the root in a simple tree topology. When the Root node sends information, all of the remaining nodes on the connection, called Leaf nodes, receive copies of the information.
- 2. Each of the Leaf Nodes on the connection can send information directly to the Root Node. The Root Node cannot distinguish which Leaf is sending information without additional (higher layer) information. (See the following note for Phase 1.)
- 3. The Leaf Nodes cannot communicate directly to each other with this connection type.

Note: Phase 1 signalling does not support traffic sent from a Leaf to the Root.

Point-to-Point Connection - a connection with only two endpoints.

Point-to-Point Protocol (PPP) - Provides a method for transmitting packets over serial point-to-point links.

Policing - the function that ensures that a network device does not accept traffic that exceeds the configured bandwidth of a connection.

Port Identifier - The identifier assigned by a logical node to represent the point of attachment of a link to that node.

Presentation Layer - Sixth layer of the OSI model, providing services to the application layer.

Primary Reference Source (PRS) - Equipment that provides a timing signal whose long-term accuracy is maintained at 1×10 -11 or better with verification to universal coordinated time (UTC) and whose timing signal is used as the basis of reference for the control of other clocks within a network.

Primitive - an abstract, implementation-independent interaction between a layer service user and a layer service provider.

Priority - the parameter of ATM connections that determines the order in which they are reduced from the peak cell rate to the sustained cell rate in times of congestion. Connections with lower priority (4 is low, 1 is high) are reduced first.

Private Branch Exchange (PBX) - a private phone system (switch) that connects to the public telephone network and offers in-house connectivity. To reach an outside line, the user must dial a digit like 8 or 9.

Private Network Node Interface or Private Network-to-Network Interface (PNNI) - a protocol that defines the interaction of private ATM switches or groups of private ATM switches

Programmable Read-Only Memory (PROM) - a chip-based information storage area that can be recorded by an operator but erased only through a physical process.

Protocol - a set of rules and formats (semantic and syntactic) that determines the communication behavior of layer entities in the performance of the layer functions.

Protocol Control Information - the information exchanged between corresponding entities using a lower layer connection to coordinate their joint operation.

Protocol Data Unit (PDU) - a unit of data specified in a layer protocol and consisting of protocol control information and layer user data.

Proxy - the process in which one system acts for another system to answer protocol requests.

Proxy Agent - an agent that queries on behalf of the manager, used to monitor objects that are not directly manageable.

Public Data Network (PDN) - a network designed primarily for data transmission and intended for sharing by many users from many organizations.

Pulse Code Modulation (PCM) - a modulation scheme that samples the information signals and transmits a series of coded pulses to represent the data.

Q.2931 - Derived from Q.93B, the narrowband ISDN signalling protocol, an ITU standard describing the signalling protocol to be used by switched virtual circuits on ATM LANs.

Quality of Service (QoS) - Quality of Service is defined on an end-to-end basis in terms of the following attributes of the end-to-end ATM connection:

Cell Loss Ratio

Cell Transfer Delay

Cell Delay Variation

Queuing Delay (QD) - refers to the delay imposed on a cell by its having to be buffered because of unavailability of resources to pass the cell onto the next network function or element. This buffering could be a result of oversubscription of a physical link, or due to a connection of higher priority or tighter service constraints getting the resource of the physical link.

Radio Frequency Interference (RFI) - the unintentional transmission of radio signals. Computer equipment and wiring can both generate and receive RFI.

Real-Time Clock - a clock that maintains the time of day, in contrast to a clock that is used to time the electrical pulses on a circuit.

Red Alarm - In T1, a red alarm is generated for a locally detected failure such as when a condition like OOF exists for 2.5 seconds, causing a CGA, (Carrier Group Alarm).

Reduced Instruction Set Computer (RISC) - a generic name for CPUs that use a simpler instruction set than more traditional designs.

Redundancy - In a data transmission, the fragments of characters and bits that can be eliminated with no loss of information.

Registration - The address registration function is the mechanism by which Clients provide address information to the LAN Emulation Server.

Relaying - a function of a layer by means of which a layer entity receives data from a corresponding entity and transmits it to another corresponding entity.

Request To Send (RTS) - an RS-232 modem interface signal (sent from the DTE to the modem on pin 4) which indicates that the DTE has data to transmit.

Requests For Comment (RFCs) - IETF documents suggesting protocols and policies of the Internet, inviting comments as to the quality and validity of those policies. These comments are collected and analyzed by the IETF in order to finalize Internet standards.

RFC1483 - Multiprotocol Encapsulation over ATM Adaptation Layer 5.

RFC1490 - Multiprotocol Interconnect over Frame Relay.

RFC1577 - Classical IP and ARP over ATM.

RFC1755 - ATM Signaling Support for IP over ATM.

Robbed-Bit Signaling - In T1, refers to the use of the least significant bit of every word of frames 6 and 12 (D4), or 6, 12, 18, and 24 (ESF) for signaling purposes.

Route Server - A physical device that runs one or more network layer routing protocols, and which uses a route query protocol in order to provide network layer routing forwarding descriptions to clients.

Router - a device that forwards traffic between networks or subnetworks based on network layer information.

Routing Domain (RD) - A group of topologically contiguous systems which are running one instance of routing.

Routing Information Protocol (RIP) - a distance vector-based protocol that provides a measure of distance, or hops, from a transmitting workstation to a receiving workstation.

Routing Protocol - A general term indicating a protocol run between routers and/or route servers in order to exchange information used to allow computation of routes. The result of the routing computation will be one or more forwarding descriptions.

SBus - hardware interface for add-in boards in later-version Sun 3 workstations.

Scalable Processor Architecture Reduced instruction set Computer (SPARC) - a powerful workstation similar to a reduced-instruction-set-computing (RISC) workstation.

Segment - a single ATM link or group of interconnected ATM links of an ATM connection.

Segmentation And Reassembly (SAR) - the SAR accepts PDUs from the CS and divides them into very small segments (44 bytes long). If the CS-PDU is less than 44 bytes, it is padded to 44 with zeroes. A two-byte header and trailer are added to this basic segment. The header identifies the message type (beginning, end, continuation, or single) and contains sequence numbering and message identification. The trailer gives the SAR-PDU payload length, exclusive of pad, and contains a CRC check to ensure the SAR-PDU integrity. The result is a 48-byte PDU that fits into the payload field of an ATM cell.

Selector (SEL) - A subfield carried in SETUP message part of ATM endpoint address Domain specific Part (DSP) defined by ISO 10589, not used for ATM network routing, used by ATM end systems only.

Semipermanent Connection - a connection established via a service order or via network management.

Serial Line IP (SLIP) - A protocol used to run IP over serial lines, such as telephone circuits or RS-232 cables, interconnecting two systems.

Service Access Point (SAP) - the point at which an entity of a layer provides services to its LM entity or to an entity of the next higher layer.

Service Data Unit (SDU) - a unit of interface information whose identity is preserved from one end of a layer connection to the other.

Service Specific Connection Oriented Protocol (SSCOP) - an adaptation layer protocol defined in ITU-T Specification: Q.2110.

Service Specific Convergence Sublayer (SSCS) - The portion of the convergence sublayer that is dependent upon the type of traffic that is being converted.

Session Layer - Layer 5 in the OSI model that is responsible for establishing and managing sessions between the application programs running in different nodes.

Severely Errored Seconds (SES) - a second during which more event errors have occurred than the SES threshold (normally 10-3).

Shaping Descriptor - *n* ordered pairs of GCRA parameters (I,L) used to define the negotiated traffic shape of an APP connection. The traffic shape refers to the load-balancing of a network, where load-balancing means configuring data flows to maximize network efficiency.

Shielded Pair - Two insulated wires in a cable wrapped with metallic braid or foil to prevent interference and provide noise free transmission.

Shielded Twisted Pair (STP) - two or more insulated wires, twisted together and then wrapped in a cable with metallic braid or foil to prevent interference and offer noise-free transmissions.

Signaling System No. 7 (SS7) - The SS7 protocol has been specified by ITU-T and is a protocol for interexchange signaling.

Simple and Efficient Adaptation Layer (SEAL) - also called AAL 5, this ATM adaptation layer assumes that higher layer processes will provide error recovery, thereby simplifying the SAR portion of the adaptation layer. Using this AAL type packs all 48 bytes of an ATM cell information field with data. It also assumes that only one message is crossing the UNI at a time. That is, multiple end-users at one location cannot interleave messages on the same VC, but must queue them for sequential transmission.

Simple Gateway Management Protocol (SGMP) - the predecessor to SNMP.

Simple Mail Transfer Protocol (SMTP) - the Internet electronic mail protocol used to transfer electronic mail between hosts.

Simple Network Management Protocol (SNMP) - the Internet standard protocol for managing nodes on an IP network.

Simple Protocol for ATM Network Signalling (SPANS) - FORE Systems' proprietary signalling protocol used for establishing SVCs between FORE Systems equipment.

Single Mode Fiber (SMF) - Fiber optic cable in which the signal or light propagates in a single mode or path. Since all light follows the same path or travels the same distance, a transmitted pulse is not dispersed and does not interfere with adjacent pulses. SMF fibers can support longer distances and are limited mainly by the amount of attenuation. Refer to MMF.

Small Computer Systems Interface (SCSI) - a standard for a controller bus that connects hardware devices to their controllers on a computer bus, typically used in small systems.

Smart PVC (SPVC) - a generic term for any communications medium which is permanently provisioned at the end points, but switched in the middle. In ATM, there are two kinds of SPVCs: smart permanent virtual path connections (SPVPCs) and smart permanent virtual channel connections (SPVCCs).

snmpd - an SMNP agent for a given adapter card.

Source - Part of communications system which transmits information.

Source Address (SA) - The address from which the message or data originated.

Source MAC Address (SA) - A six octet value uniquely identifying an end point and which is sent in an IEEE LAN frame header to indicate source of frame.

Source Traffic Descriptor - a set of traffic parameters belonging to the ATM Traffic Descriptor used during the connection set-up to capture the intrinsic traffic characteristics of the connection requested by the source.

Spanning Tree Protocol - provides loop-free topology in a network environment where there are redundant paths.

Static Route - a route that is entered manually into the routing table.

Statistical Multiplexing - a technique for allowing multiple channels and paths to share the same link, typified by the ability to give the bandwidth of a temporarily idle channel to another channel.

Stick and Click (SC) - Designation for an Optical Connector featuring a 2.5 mm physically contacting ferrule with a push-pull mating design. Commonly referred to as Structured Cabling, Structured Connectors or Stick and Click

Stick and Turn (ST) - A fiber-optic connector designed by AT&T which uses the bayonet style coupling rather than screw-on as the SMA uses. The ST is generally considered the eventual replacement for the SMA type connector.

Store-and-Forward - the technique of receiving a message, storing it until the proper outgoing line is available, then retransmitting it, with no direct connection between incoming and outgoing lines.

Straight Tip (ST) - see Stick and Turn.

Structured Cabling (SC) - see Stick and Click.

Structured Connectors (SC) - see Stick and Click.

Sublayer - a logical subdivision of a layer.

SubNetwork Access Protocol (SNAP) - a specially reserved variant of IEEE 802.2 encoding SNAP indicates to look further into the packet where it will fine a Type field.

Subscriber Network Interface (SNI) - the interface between an SMDS end user's CPE and the network directly serving the end user, supported by either a DS1 or DS3 access arrangement.

Super Frame (SF) - a term used to describe the repeating 12 D4 frame format that composes a standard (non-ESF) T1 service.

Super User - a login ID that allows unlimited access to the full range of a device's functionality, including especially the ability to reconfigure the device and set passwords.

Sustainable Cell Rate (SCR) - ATM Forum parameter defined for traffic management. For VBR connections, SCR determines the long-term average cell rate that can be transmitted.

Sustained Information Rate (SIR) - In ATM this refers to the long-term average data transmission rate across the User-to-Network Interface. In SMDS this refers to the committed information rate (similar to CIR for Frame Relay Service).

Switch - Equipment used to interconnect lines and trunks.

Switched Connection - A connection established via signaling.

Switched Multimegabit Data Service (SMDS) - a high-speed, datagram-based, public data network service expected to be widely used by telephone companies in their data networks.

Switched Virtual Channel Connection (SVCC) - A Switched VCC is one which is established and taken down dynamically through control signaling. A Virtual Channel Connection (VCC) is an ATM connection where switching is performed on the VPI/VCI fields of each cell.

Switched Virtual Circuit (or Channel) (SVC) - a channel established on demand by network signalling, used for information transport between two locations and lasting only for the duration of the transfer; the datacom equivalent of a dialed telephone call.

Switched Virtual Path Connection (SVPC) - a connection which is established and taken down dynamically through control signaling. A Virtual Path Connection (VPC) is an ATM connection where switching is performed on the VPI field only of each cell.

Switching System - A set of one or more systems that act together and appear as a single switch for the purposes of PNNI routing.

Symmetric Connection - a connection with the same bandwidth specified for both directions.

Synchronous - signals that are sourced from the same timing reference and hence are identical in frequency.

Synchronous Data Link Control (SDLC) - IBM's data link protocol used in SNA networks.

Synchronous Optical Network (SONET) - a body of standards that defines all aspects of transporting and managing digital traffic over optical facilities in the public network.

Synchronous Payload Envelope (SPE) - the payload field plus a little overhead of a basic SONET signal.

Synchronous Transfer Mode (STM) - a transport and switching method that depends on information occurring in regular, fixed patterns with respect to a reference such as a frame pattern.

Synchronous Transport Signal (STS) - a SONET electrical signal rate.

Systeme En Coleur Avec Memoire (SECAM) - Sequential and Memory Color Television - Started in France in the late 60s, and used by other countries with a political affiliation. This is. The B-Y and R-Y signals are transmitted on alternate lines modulated on an FM subcarrier. The memory is a one line delay line in the receiver to make both color difference signals available at the same time on all lines. Due to FM, the signal is robust in difficult terrain.

Systems Network Architecture (SNA) - a proprietary networking architecture used by IBM and IBM-compatible mainframe computers.

T1 - a specification for a transmission line. The specification details the input and output characteristics and the bandwidth. T1 lines run at 1.544 Mbps and provide for 24 data channels. In common usage, the term "T1" is used interchangeably with "DS1."

T1 Link - A wideband digital carrier facility used for transmission of digitized voice, digital data, and digitized image traffic. This link is composed of two twisted-wire pairs that can carry 24 digital channels, each operating at 64K bps at the aggregate rate of 1.544M bps, full duplex. Also referred to as DS-1.

T3 - a specification for a transmission line, the equivalent of 28 T1 lines. T3 lines run at 44.736 Mbps. In common usage, the term "T3" is used interchangeably with "DS3."

Tachometer - in *ForeView*, the tachometer shows the level of activity on a given port. The number in the tachometer shows the value of a chosen parameter in percentage, with a colored bar providing a semi-logarithmic representation of that percentage.

Tagged Cell Rate (TCR) - An ABR service parameter, TCR limits the rate at which a source may send out-of-rate forward RM-cells. TCR is a constant fixed at 10 cells/second.

Telephony - The conversion of voices and other sounds into electrical signals which are then transmitted by telecommunications media.

Telnet - a TCP/IP protocol that defines a client/server mechanism for emulating directly-connected terminal connections.

Terminal Equipment (TE) - Terminal equipment represents the endpoint of ATM connection(s) and termination of the various protocols within the connection(s).

Throughput - Measurement of the total useful information processed or communicated by a computer during a specified time period, i.e. packets per second.

Time Division Multiplexing (TDM) - a method of traditional digital multiplexing in which a signal occupies a fixed, repetitive time slot within a higher-rate signal.

Token Ring - a network access method in which the stations circulate a token. Stations with data to send must have the token to transmit their data.

topology - a program that displays the topology of a FORE Systems ATM network. An updated topology can be periodically re-displayed by use of the interval command option.

Traffic - the calls being sent and received over a communications network. Also, the packets that are sent on a data network.

Traffic Management (TM) - The traffic control and congestion control procedures for ATM. ATM layer traffic control refers to the set of actions taken by the network to avoid congestion conditions. ATM layer congestion control refers to the set of actions taken by the network to minimize the intensity, spread and duration of congestion. The following functions form a framework for managing and controlling traffic and congestion in ATM networks and may be used in appropriate combinations:

Connection Admission Control Feedback Control Usage Parameter Control Priority Control Traffic Shaping Network Resource Management Frame Discard ABR Flow Control

Traffic Parameter - A parameter for specifying a particular traffic aspect of a connection.

Trailer - the protocol control information located at the end of a PDU.

Transit Delay - the time difference between the instant at which the first bit of a PDU crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.

Transmission Control Protocol (TCP) - a specification for software that bundles and unbundles sent and received data into packets, manages the transmission of packets on a network, and checks for errors.

Transmission Control Protocol/Internet Protocol (TCP/IP) - a set of communications protocols that has evolved since the late 1970s, when it was first developed by the Department of Defense. Because programs supporting these protocols are available on so many different computer systems, they have become an excellent way to connect different types of computers over networks.

Transmission Convergence (TC) - generates and receives transmission frames and is responsible for all overhead associated with the transmission frame. The TC sublayer packages cells into the transmission frame.

Transmission Convergence Sublayer (TCS) - This is part of the ATM physical layer that defines how cells will be transmitted by the actual physical layer.

Transparent Asynchronous Transmitter/Receiver Interface (TAXI) - Encoding scheme used for FDDI LANs as well as for ATM; supports speed typical of 100 Mbps over multimode fiber.

Transport Layer - Layer Four of the OSI reference model that is responsible for maintaining reliable end-to-end communications across the network.

trap - a program interrupt mechanism that automatically updates the state of the network to remote network management hosts. The SNMP agent on the switch supports these SNMP traps.

Trivial File Transfer Protocol (TFTP) - Part of IP, a simplified version of FTP that allows files to be transferred from one computer to another over a network.

Twisted Pair - Insulated wire in which pairs are twisted together. Commonly used for telephone connections, and LANs because it is inexpensive.

Unassigned Cells - a generated cell identified by a standardized virtual path identifier (VPI) and virtual channel identifier (VCI) value, which does not carry information from an application using the ATM Layer service.

Unavailable Seconds (UAS) - a measurement of signal quality. Unavailable seconds start accruing when ten consecutive severely errored seconds occur.

UNI 3.0/3.1 - the User-to-Network Interface standard set forth by the ATM Forum that defines how private customer premise equipment interacts with private ATM switches.

Unicasting - The transmit operation of a single PDU by a source interface where the PDU reaches a single destination.

Universal Test & Operations Interface for ATM (UTOPIA) - Refers to an electrical interface between the TC and PMD sublayers of the PHY layer.

Unshielded Twisted Pair (UTP) - a cable that consists of two or more insulated conductors in which each pair of conductors are twisted around each other. There is no external protection and noise resistance comes solely from the twists.

Unspecified Bit Rate (UBR) - a type of traffic that is not considered time-critical (e.g., ARP messages, pure data), allocated whatever bandwidth is available at any given time. UBR traffic is given a "best effort" priority in an ATM network with no guarantee of successful transmission.

Uplink - Represents the connectivity from a border node to an upnode.

Usage Parameter Control (UPC) - mechanism that ensures that traffic on a given connection does not exceed the contracted bandwidth of the connection, responsible for policing or enforcement. UPC is sometimes confused with congestion management (see *congestion management*).

User Datagram Protocol (UDP) - the TCP/IP transaction protocol used for applications such as remote network management and name-service access; this lets users assign a name, such as "RVAX*2,S," to a physical or numbered address.

User-to-Network Interface (UNI) - the physical and electrical demarcation point between the user and the public network service provider.

V.35 - ITU-T standard describing a synchronous, physical layer protocol used for communications between a network access device and a packet network. V.35 is most commonly used in the United States and Europe, and is recommended for speeds up to 48 Kbps.

Variable Bit Rate (VBR) - a type of traffic that, when sent over a network, is tolerant of delays and changes in the amount of bandwidth it is allocated (e.g., data applications).

Virtual Channel (or Circuit) (VC) - a communications path between two nodes identified by label rather than fixed physical path.

Virtual Channel Connection (VCC) - a unidirectional concatenation of VCLs that extends between the points where the ATM service users access the ATM Layer. The points at which the ATM cell payload is passed to, or received from, the users of the ATM Layer (i.e., a higher layer or ATMM-entity) for processing signify the endpoints of a VCC.

Virtual Channel Identifier (VCI) - the address or label of a VC; a value stored in a field in the ATM cell header that identifies an individual virtual channel to which the cell belongs. VCI values may be different for each data link hop of an ATM virtual connection.

Virtual Channel Link (VCL) - a means of unidirectional transport of ATM cells between the point where a VCI value is assigned and the point where that value is translated or removed.

Virtual Channel Switch - a network element that connects VCLs. It terminates VPCs and translates VCI values. The Virtual Channel Switch is directed by Control Plane functions and relays the cells of a VC.

Virtual Connection - an endpoint-to-endpoint connection in an ATM network. A virtual connection can be either a virtual path or a virtual channel.

Virtual Local Area Network (VLAN) - Work stations connected to an intelligent device which provides the capabilities to define LAN membership.

Virtual Network Software (VINES) - Banyan's network operating system based on UNIX and its protocols.

Virtual Path (VP) - a unidirectional logical association or bundle of VCs.

Virtual Path Connection (VPC) - a concatenation of VPLs between virtual path terminators (VPTs). VPCs are unidirectional.

Virtual Path Identifier (VPI) - the address or label of a particular VP; a value stored in a field in the ATM cell header that identifies an individual virtual path to which the cell belongs. A virtual path may comprise multiple virtual channels.

Virtual Path Link (VPL) - a means of unidirectional transport of ATM cells between the point where a VPI value is assigned and the point where that value is translated or removed.

Virtual Path Switch - a network element that connects VPLs, it translates VPI (not VCI) values and is directed by Control Plane functions. The Virtual Path Switch relays the cells of a Virtual Path.

Virtual Path Terminator (VPT) - a system that unbundles the VCs of a VP for independent processing of each VC.

Virtual Private Data Network (VPDN) - a private data communications network built on public switching and transport facilities rather than dedicated leased facilities such as T1s.

Virtual Private Network (VPN) - a private voice communications network built on public switching and transport facilities rather than dedicated leased facilities such as T1s.

Virtual Source/Virtual Destination (VS/VD) - An ABR connection may be divided into two or more separately controlled ABR segments. Each ABR control segment, except the first, is sourced by a virtual source. A virtual source implements the behavior of an ABR source endpoint. Backwards RM-cells received by a virtual source are removed from the connection. Each ABR control segment, except the last, is terminated by a virtual destination. A virtual destination assumes the behavior of an ABR destination endpoint. Forward RM-cells received by a virtual destination are turned around and not forwarded to the next segment of the connection.

Virtual Tributary (VT) - a structure used to carry payloads such as DS1s that run at significantly lower rates than STS-1s.

Warm Start Trap - an SNMP trap which indicates that SNMP alarm messages or agents have been enabled.

Wide-Area Network (WAN) - a network that covers a large geographic area.

Wideband Channel - Communications channel with more capacity (19.2K bps) than the standard capacity of a voice grade line.

X.21 - ITU-T standard for serial communications over synchronous digital lines. The X.21 protocol is used primarily in Europe and Japan.

X.25 - a well-established data switching and transport method that relies on a significant amount of processing to ensure reliable transport over metallic media.

Yellow Alarm - An alarm signal sent back toward the source of a failed signal due to the presence of an AIS (may be used by APS equipment to initiate switching).

Zero Byte Time Slot Interchange (ZBTSI) - A technique used with the T carrier extended superframe format (ESF) in which an area in the ESF frame carries information about the location of all-zero bytes (eight consecutive "0"s) within the data stream.

Zero Code Suppression - The insertion of a "1" bit to prevent the transmission of eight or more consecutive "0" bits. Used primarily with T1 and related digital telephone company facilities, which require a minimum "1's density" in order to keep the individual subchannels of a multiplexed, high speed facility active.

Zero-Bit Insertion - A technique used to achieve transparency in bit-oriented protocols. A zero is inserted into sequences of one bits that cause false flag direction.

Index

A	Circuit Emulation Services (CES) 4 - 1
ABR 1 - 31	Classical IP ATM network
address registration	configuring 2 - 1
ARP cache 2 - 11	Classical IP interface 2 - 2
ARP reply 2 - 11	Classical IP over ATM 2 - 1
ARP request 2 - 11	Classical IP PVC 2 - 13
ARP server	CLP bit
configuring a FORE switch as 2 - 10	connection
configuring a Hewlett-Packard	configuration-direct
computer as 2 - 9	control-direct 3 - 4
configuring a host as 2 - 7	control-distribute 3 - 4
one time 2 - 8	data-direct
configuring a host as	multicast-forward 3 - 4
permanently2 - 9	multicast-send 3 - 4
ATM ARP (ATM address resolution	Constant Bit Rate (CBR) 1 - 31
protocol) 2 - 6	D
В	defining an ELAN 3 - 12
Broadcast and Unknown Server (BUS) 3 - 3	dynamic path
BUS 3 - 3	•
starting 3 - 23	E
С	ELAN 3 - 1
CBR	configuring
CDVT 1 - 32	joining 3 - 24
Cell Delay Variation Tolerance	Emulated LAN 3 - 1
(CDVT)	address resolution
Cell Loss Priority (CLP) bit 1 - 33	components 3 - 2
CES (Circuit Emulation Services) 4 - 1	data transfer 3 - 7
CES connection	initialization
creating	operation
displaying 4 - 4	registration 3 - 6
	ESI (End System Identifier) 2 - 4

F	LECS 3 - 3
FORE IP	well-known address 3 - 24
ForeThought PNNI B - 1	LECS configuration file
backbone topology B - 9	default
border switches B - 8	sample
hello indication messages B - 8	starting 3 - 21
link metrics	syntax 3 - 10
peer group summary node B - 8, B - 9	LECS control parameters 3 - 15
physical network B - 5	LES
topology database B - 5	starting 3 - 23
two-level hierarchy B - 6	link-scope UNI C - 7
•	LIS (Logical IP Subnet) 2 - 2
G GCRA	Logical IP Subnet (LIS) 2 - 2
	logical link (loglink) B - 2
Generic Cell Rate Algorithm (GCRA) 1 - 32	loglink (logical link) B - 2
I	M
ILMI (Interim Local Management	
Interface) 2 - 4	Maximum Burst Size (MBS) 1 - 31
InARP reply 2 - 11	MBS 1 - 31
InARP request 2 - 11	N
Interim Local Management Interface	network configuration examples 2 - 15
(ILMI)	NSAP addresses
inverse ARP (InARP) 2 - 11	NSAP prefix 2 - 4
L	0
LAN Emulation Client (LEC) 3 - 3	originating path 1 - 6
LAN Emulation Configuration Server	originating paths
(LECS)	adding 1 - 21
LAN Emulation Server (LES) 3 - 3	displaying advanced information
LAN Emulation Services	for all 1 - 11
LAN Emulation services	displaying all 1 - 9
starting 3 - 21	
leaky bucket algorithm 1 - 32	P
LEC	PCR
failover mechanism 3 - 17	Peak Cell Rate (PCR)
starting 3 - 24	PNNI B - 1
LEC failover mechanism 3 - 26	PVC revalidation 2 - 13

R
RFC-1577 2 - 1
S
SCR 1 - 31, 1 - 32
selector field 2 - 4
SNMP indexing
SNMP trap
adding
deleting
displaying
supported on the switch $\dots A$ - 3
SPANS 1 - 27
SPANS interface 2 - 3
SPANS SPVC
creating 1 - 27
SPANS SPVCs
displaying 1 - 30
SRTS (Synchronous Residual Time
Stamp)
Sustainable Cell Rate (SCR) 1 - 31, 1 - 32
Synchronous Residual Time Stamp (SRTS)
T
tagging cells
Technical Supportii
terminating path
terminating paths
adding
displaying advanced information for all 1 - 11
displaying all
through paths
displaying advanced information
for all
traffic policing

U
UBR1 - 31
Unspecified Bit Rate (UBR) 1 - 31
UPC1 - 32
UPC contract
adding 1 - 35
UPC traffic contract parameters $\dots 1$ - 33
Usage Parameter Control (UPC) 1 - 32
V
• Variable Bit Rate (VBR) 1 - 31
VBR
VCI allocation range
VC-space
victual channel
virtual channel identifier (VCI) 1 - 1
virtual path
displaying advanced information
for all
virtual path identifier (VPI)
virtual path terminator
displaying advanced information
for all 1 - 10
virtual path terminators $\ \dots \ 1$ - 6
VP-scope UNI

Index